

Data Replication Service

Real-Time Synchronization

Issue 54
Date 2024-11-30



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <https://e.huawei.com>

Security Declaration

Product Lifecycle

Huawei's regulations on product lifecycle are subject to the *Product End of Life Policy*. For details about this policy, visit the following web page:

<https://support.huawei.com/ecolumnsweb/en/warranty-policy>

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Initial Digital Certificate

The Initial digital certificates on Huawei devices are subject to the *Rights and Responsibilities of Initial Digital Certificates on Huawei Devices*. For details about this document, visit the following web page:

<https://support.huawei.com/enterprise/en/bulletins-service/ENEWS2000015789>

Huawei Enterprise End User License Agreement

This agreement is the end user license agreement between you (an individual, company, or any other entity) and Huawei for the use of the Huawei Software. Your use of the Huawei Software will be deemed as your acceptance of the terms mentioned in this agreement. For details about this agreement, visit the following web page:

<https://e.huawei.com/en/about/eula>

Lifecycle of Product Documentation

Huawei after-sales user documentation is subject to the *Product Documentation Lifecycle Policy*. For details about this policy, visit the following web page:

<https://support.huawei.com/enterprise/en/bulletins-website/ENEWS2000017761>

Contents

1 Synchronization Overview.....	1
2 Data Synchronization Topologies.....	19
3 To the Cloud.....	22
3.1 From MySQL to MySQL.....	22
3.2 From MySQL to PostgreSQL.....	45
3.3 From MySQL to GaussDB Distributed.....	62
3.4 From MySQL to GaussDB Primary/Standby.....	84
3.5 From MySQL to GaussDB(DWS).....	104
3.6 From MySQL to GaussDB(for MySQL).....	126
3.7 From MySQL to MariaDB.....	146
3.8 From PostgreSQL to PostgreSQL.....	167
3.9 From PostgreSQL to GaussDB(DWS).....	191
3.10 From PostgreSQL to GaussDB Primary/Standby.....	211
3.11 From PostgreSQL to GaussDB Distributed.....	231
3.12 From Oracle to MySQL.....	252
3.13 From Oracle to GaussDB(for MySQL).....	278
3.14 From Oracle to GaussDB Primary/Standby.....	300
3.15 From Oracle to GaussDB Distributed.....	326
3.16 From Oracle to DDM.....	351
3.17 From Oracle to GaussDB(DWS).....	371
3.18 From Oracle to PostgreSQL.....	392
3.19 From DDM to MySQL.....	412
3.20 From DDM to GaussDB(DWS).....	429
3.21 From DDM to DDM.....	444
3.22 From DB2 for LUW to GaussDB Primary/Standby.....	458
3.23 From DB2 for LUW to GaussDB Distributed.....	472
3.24 From DB2 for LUW to GaussDB(DWS).....	486
3.25 From TiDB to GaussDB(for MySQL).....	501
3.26 From Microsoft SQL Server to GaussDB(DWS).....	516
3.27 From Microsoft SQL Server to GaussDB Primary/Standby.....	533
3.28 From Microsoft SQL Server to GaussDB Distributed.....	550
3.29 From Microsoft SQL Server to Microsoft SQL Server.....	567

3.30 From Microsoft SQL Server to MySQL.....	584
3.31 From Microsoft SQL Server to GaussDB(for MySQL).....	603
3.32 From Microsoft SQL Server to PostgreSQL.....	622
3.33 From MongoDB to DDS.....	641
3.34 From MariaDB to MariaDB.....	656
3.35 From MariaDB to MySQL.....	674
3.36 From MariaDB to GaussDB(for MySQL).....	694
3.37 From GaussDB(for MySQL) to GaussDB(for MySQL).....	713
3.38 From Cassandra to GeminiDB Cassandra.....	733
3.39 From Dynamo to GeminiDB Cassandra.....	746
3.40 From Dynamo to GeminiDB Dynamo.....	759
3.41 From AWS DocumentDB to DDS.....	771
3.42 From Redis to GeminiDB Redis.....	786
3.43 From Redis Cluster to GeminiDB Redis.....	799
4 Out of the Cloud.....	812
4.1 From MySQL to MySQL.....	812
4.2 From MySQL to Kafka.....	835
4.3 From MySQL to CSS/ES.....	854
4.4 From MySQL to Oracle.....	872
4.5 From MySQL to MariaDB.....	889
4.6 From DDM to MySQL.....	910
4.7 From DDM to Oracle.....	926
4.8 From DDM to Kafka.....	941
4.9 From DDS to MongoDB.....	953
4.10 From DDS to Kafka.....	967
4.11 From PostgreSQL to PostgreSQL.....	978
4.12 From PostgreSQL to Kafka.....	1001
4.13 From GaussDB Distributed to MySQL.....	1015
4.14 From GaussDB Distributed to Oracle.....	1036
4.15 From GaussDB Distributed to GaussDB(DWS).....	1057
4.16 From GaussDB Distributed to Kafka.....	1077
4.17 From GaussDB Distributed to GaussDB Distributed.....	1094
4.18 From GaussDB Distributed to GaussDB Primary/Standby.....	1117
4.19 From GaussDB Distributed to PostgreSQL.....	1139
4.20 From GaussDB Primary/Standby to MySQL.....	1158
4.21 From GaussDB Primary/Standby to Oracle.....	1179
4.22 From GaussDB Primary/Standby to GaussDB(DWS).....	1200
4.23 From GaussDB Primary/Standby to Kafka.....	1220
4.24 From GaussDB Primary/Standby to GaussDB Distributed.....	1236
4.25 From GaussDB Primary/Standby to GaussDB Primary/Standby.....	1258
4.26 From GaussDB Primary/Standby to PostgreSQL.....	1280
4.27 From GaussDB(for MySQL) to MySQL.....	1299

4.28 From GaussDB(for MySQL) to GaussDB(DWS).....	1318
4.29 From GaussDB(for MySQL) to Kafka.....	1336
4.30 From GaussDB(for MySQL) to CSS/ES.....	1350
4.31 From GaussDB(for MySQL) to Oracle.....	1368
4.32 From MariaDB to MariaDB.....	1382
4.33 From Microsoft SQL Server to Kafka.....	1399
4.34 From GeminiDB Redis to Redis.....	1412
4.35 From GeminiDB Redis to Redis Cluster.....	1426
5 Between Self-built Databases.....	1439
5.1 From MySQL to Kafka.....	1439
5.2 From MySQL to CSS/ES.....	1457
5.3 From MySQL to GaussDB Primary/Standby.....	1475
5.4 From MySQL to GaussDB Distributed.....	1494
5.5 From Oracle to Kafka.....	1514
5.6 From Oracle to GaussDB Primary/Standby.....	1531
5.7 From Oracle to GaussDB Distributed.....	1557
5.8 From PostgreSQL to Kafka.....	1582
5.9 From PostgreSQL to GaussDB Primary/Standby.....	1596
5.10 From PostgreSQL to GaussDB Distributed.....	1615
5.11 From GaussDB Primary/Standby to MySQL.....	1634
5.12 From GaussDB Primary/Standby to Oracle.....	1656
5.13 From GaussDB Primary/Standby to Kafka.....	1677
5.14 From GaussDB Primary/Standby to GaussDB Distributed.....	1692
5.15 From GaussDB Primary/Standby to GaussDB Primary/Standby.....	1715
5.16 From GaussDB Primary/Standby to PostgreSQL.....	1736
5.17 From GaussDB Distributed to MySQL.....	1757
5.18 From GaussDB Distributed to Oracle.....	1780
5.19 From GaussDB Distributed to Kafka.....	1801
5.20 From GaussDB Distributed to GaussDB Distributed.....	1818
5.21 From GaussDB Distributed to GaussDB Primary/Standby.....	1841
5.22 From GaussDB Distributed to PostgreSQL.....	1864
5.23 From DB2 for LUW to GaussDB Primary/Standby.....	1884
5.24 From DB2 for LUW to GaussDB Distributed.....	1898
5.25 From Microsoft SQL Server to Kafka.....	1912
6 Two-Way Synchronization.....	1925
6.1 Two-Way Synchronization from MySQL to MySQL.....	1925
7 Task Management.....	1946
7.1 Creating a Synchronization Task.....	1946
7.2 Querying the Synchronization Progress.....	1963
7.3 Viewing Synchronization Logs.....	1967
7.4 Data Comparison (Comparing Synchronization Items).....	1967

7.5 Managing Objects.....	1988
7.5.1 Editing Synchronization Objects.....	1988
7.5.2 Importing Synchronization Objects.....	1992
7.5.3 Changing Object Names (Mapping Object Names).....	1993
7.5.4 Viewing Synchronization Mapping Information.....	1999
7.5.5 Processing Data.....	1999
7.6 Managing Parameters.....	2012
7.6.1 Changing Task Parameters.....	2012
7.6.2 Viewing Parameter Change History.....	2046
7.7 Task Life Cycle.....	2047
7.7.1 Viewing Task Details.....	2047
7.7.2 Modifying Task Information.....	2047
7.7.3 Modifying Connection Information.....	2049
7.7.4 Modifying the Flow Control Mode.....	2050
7.7.5 Changing the Synchronization Mode.....	2054
7.7.6 Editing a Synchronization Task.....	2054
7.7.7 Resuming a Synchronization Task.....	2055
7.7.8 Skipping Data and Resuming a Synchronization Task.....	2056
7.7.9 Pausing a Synchronization Task.....	2057
7.7.10 Resetting a Synchronization Task.....	2061
7.7.11 Resetting Synchronization Position.....	2064
7.7.12 Restarting a Synchronization Task.....	2064
7.7.13 Skipping DDL Statements.....	2065
7.7.14 Performing a Switchover for a Dual-AZ Task.....	2066
7.7.15 Exchanging the Direction for a Two-Way Synchronization Task.....	2067
7.7.16 Changing a Single-AZ Task to a Dual-AZ Task.....	2068
7.7.17 Cloning a Synchronization Task.....	2069
7.7.18 Changing Specifications.....	2070
7.7.19 Unsubscribing from a Yearly/Monthly Task.....	2071
7.7.20 Stopping a Synchronization Task.....	2073
7.7.21 Deleting a Synchronization Task.....	2074
7.7.22 Importing Tasks in Batches.....	2075
7.7.23 Filtering Large Fields.....	2076
7.7.24 Upgrading the Version of a DRS Task.....	2078
7.7.25 Task Statuses.....	2080
8 Tag Management.....	2083
9 Connection Diagnosis.....	2085
10 Viewing Abnormal Data.....	2087
11 Interconnecting with CTS.....	2088
11.1 Key Operations Recorded by CTS.....	2088
11.2 Viewing Traces.....	2088

12 Interconnecting with Cloud Eye.....	2090
12.1 Supported Metrics.....	2090
12.2 Configuring Alarm Rules.....	2095
12.3 Viewing Monitoring Metrics.....	2096
13 Interconnecting with LTS.....	2098
13.1 Log Reporting.....	2098
13.2 Viewing and Downloading Logs.....	2099
14 Operation Reference in Synchronization Scenarios.....	2102
14.1 Kafka Message Format.....	2102
14.2 Kafka Authentication.....	2111
14.3 Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database.....	2114
14.4 Forcibly Stopping Synchronization of GaussDB Distributed.....	2115
14.5 Forcibly Stopping Synchronization of GaussDB Primary/Standby.....	2116
14.6 Forcibly Stopping Synchronization of PostgreSQL.....	2117
14.7 Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.....	2119
15 Appendix.....	2123
15.1 General Usage Suggestions.....	2123
15.2 DRS Pre-Check Items.....	2124

1 Synchronization Overview

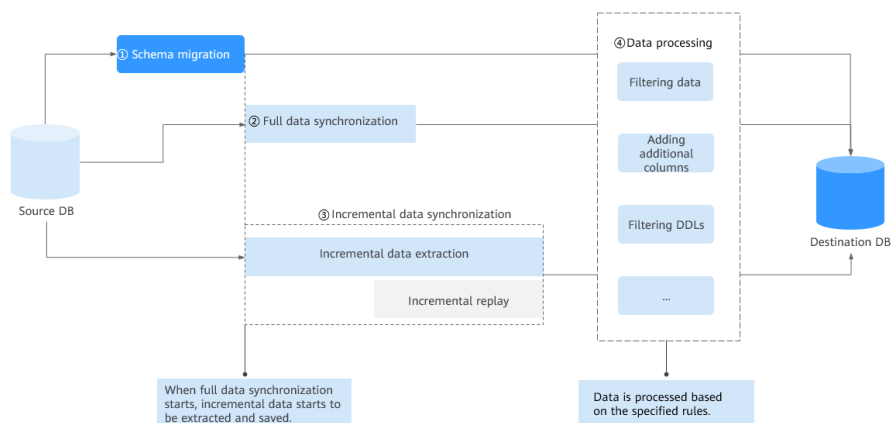
Real-time synchronization refers to the real-time flow of key service data from sources to destinations while consistency of data can be ensured.

It is different from migration. Migration means moving your overall database from one platform to another. Synchronization refers to the continuous flow of data between different services.

You can use real-time synchronization in many scenarios such as real-time analysis, report system, and data warehouse environment.

Real-time synchronization is mainly used for synchronizing tables and data. It can meet various requirements, such as many-to-one, one-to-many synchronization, dynamic addition and deletion of tables, and synchronization between tables with different names.

Figure 1-1 Real-time synchronization principle



Supported Database Types

The following table lists the source database and destination database types supported by DRS in real-time synchronization.

Table 1-1 Synchronization scheme

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
<ul style="list-style-type: none"> • On-premises MySQL databases • MySQL databases on an ECS • MySQL databases on other clouds • RDS for MySQL 	RDS for MySQL	Incremental Full Full+Incremental	From MySQL to MySQL (To the cloud)
	GaussDB(for MySQL)	Incremental Full+Incremental	From MySQL to GaussDB(for MySQL)
	RDS for PostgreSQL	Full Full+Incremental	From MySQL to PostgreSQL
	GaussDB distributed	Incremental Full Full+Incremental	From MySQL to GaussDB Distributed (To the cloud) From MySQL to GaussDB Distributed (self-built - self-built)
	GaussDB primary/standby	Incremental Full Full+Incremental	From MySQL to GaussDB Primary/Standby (To the cloud) From MySQL to GaussDB Primary/Standby (self-built - self-built)
	GaussDB(DWS) NOTE This service is currently not supported. Use DataArts Studio for data synchronization.	Incremental Full Full+Incremental	From MySQL to GaussDB(DWS) DataArts Studio
	RDS for MariaDB	Incremental Full Full+Incremental	From MySQL to MariaDB (To the cloud)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
RDS for MySQL	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds 	Incremental Full+Incremental	From MySQL to MySQL (Out of the cloud)
	Kafka	Incremental Full+Incremental	From MySQL to Kafka (Out of the cloud)
	CSS/ES	Full+Incremental	From MySQL to CSS/ES (Out of the cloud)
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	Full+Incremental	From MySQL to Oracle
	<ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds 	Incremental Full Full+Incremental	From MySQL to MariaDB (Out of the cloud)
<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS 	Kafka	Incremental Full+Incremental	From MySQL to Kafka (self-built - self-built)
	CSS/ES	Full+Incremental	From MySQL to CSS/ES (self-built - self-built)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
<ul style="list-style-type: none"> On-premises PostgreSQL databases PostgreSQL databases on an ECS PostgreSQL databases on other clouds RDS for PostgreSQL 	RDS for PostgreSQL	Incremental Full Full+Incremental	From PostgreSQL to PostgreSQL
	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full+Incremental	From PostgreSQL to GaussDB(DWS)
	GaussDB primary/standby	Incremental Full Full+Incremental NOTE Only whitelisted users can perform the incremental data synchronization.	From PostgreSQL to GaussDB Primary/Standby
	GaussDB distributed	Incremental Full Full+Incremental NOTE Only whitelisted users can perform the incremental data synchronization.	From PostgreSQL to GaussDB Distributed
<ul style="list-style-type: none"> On-premises PostgreSQL databases PostgreSQL databases on an ECS 	Kafka	Incremental	From PostgreSQL to Kafka (self-built - self-built)
RDS for PostgreSQL	<ul style="list-style-type: none"> On-premises PostgreSQL databases PostgreSQL databases on ECSs 	Incremental Full Full+Incremental	From PostgreSQL to PostgreSQL (Out of the cloud)
	Kafka	Incremental	From PostgreSQL to Kafka (Out of the cloud)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
DDM	RDS for MySQL	Incremental Full Full+Incremental	From DDM to MySQL (To the cloud)
	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds 	Incremental Full Full+Incremental	From DDM to MySQL (Out of the cloud)
	DDM	Full+Incremental	From DDM to DDM
	GaussDB(DWS) NOTE This service is currently not supported. Use DataArts Studio for data synchronization.	Full+Incremental	From DDM to GaussDB(DWS) DataArts Studio
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	Incremental Full Full+Incremental	From DDM to Oracle
	Kafka	Incremental	From DDM to Kafka
<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	RDS for MySQL	Incremental Full Full+Incremental	From Oracle to MySQL
	GaussDB(for MySQL)	Full Full+Incremental	From Oracle to GaussDB(for MySQL)
	RDS for PostgreSQL	Full Full+Incremental	From Oracle to PostgreSQL

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
	GaussDB primary/standby	Incremental Full Full+Incremental	From Oracle to GaussDB Primary/Standby (To the cloud) From Oracle to GaussDB Primary/Standby (self-built - self-built)
	GaussDB distributed	Incremental Full Full+Incremental	From Oracle to GaussDB Distributed (To the cloud) From Oracle to GaussDB Distributed (self-built - self-built)
	DDM	Full Full+Incremental	From Oracle to DDM
	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Incremental Full Full+Incremental	From Oracle to GaussDB(DWS)
	Kafka	Incremental	From Oracle to Kafka
GaussDB distributed	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL 	Full Incremental Full+Incremental	From GaussDB Distributed to MySQL (Out of the cloud) From GaussDB Distributed to MySQL (self-built - self-built)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	Full Incremental Full+Incremental	From GaussDB Distributed to Oracle (Out of the cloud) From GaussDB Distributed to Oracle (self-built - self-built)
	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full Incremental Full+Incremental	From GaussDB Distributed to GaussDB(DWS)
	Kafka	Incremental	From GaussDB Distributed to Kafka (Out of the cloud) From GaussDB Distributed to Kafka (self-built - self-built)
	GaussDB distributed	Full Incremental Full+Incremental	From GaussDB Distributed to GaussDB Distributed From GaussDB Distributed to GaussDB Distributed (self-built - self-built)
	GaussDB primary/standby	Full Incremental Full+Incremental	From GaussDB Distributed to GaussDB Primary/Standby (Out of the cloud) From GaussDB Distributed to GaussDB Primary/Standby (self-built - self-built)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
GaussDB primary/standby	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL 	Full Incremental Full+Incremental	From GaussDB Primary/Standby to MySQL (Out of the cloud) From GaussDB Primary/Standby to MySQL (self-built - self-built)
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	Full Incremental Full+Incremental	From GaussDB Primary/Standby to Oracle (Out of the cloud) From GaussDB Primary/Standby to Oracle (self-built - self-built)
	Kafka	Incremental	From GaussDB Primary/Standby to Kafka (Out of the cloud) From GaussDB Primary/Standby to Kafka (self-built - self-built)
	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB(DWS)
	GaussDB distributed	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB Distributed (Out of the cloud) From GaussDB Primary/Standby to GaussDB Distributed (self-built - self-built)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
	GaussDB primary/standby	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB Primary/Standby From GaussDB Primary/Standby to GaussDB Primary/Standby (self-built - self-built)
GaussDB(for MySQL)	GaussDB(for MySQL)	Incremental Full+Incremental	From GaussDB(for MySQL) to GaussDB(for MySQL)
	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL 	Incremental Full+Incremental	From GaussDB(for MySQL) to MySQL
	GaussDB(DWS) NOTE This service is currently not supported. Use DataArts Studio for data synchronization.	Full+Incremental	From GaussDB(for MySQL) to GaussDB(DWS) DataArts Studio
	Kafka	Incremental Full+Incremental	From GaussDB(for MySQL) to Kafka
	CSS/ES	Full+Incremental	From GaussDB(for MySQL) to CSS/ES

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	Full+Incremental	From GaussDB(for MySQL) to Oracle
DDS	<ul style="list-style-type: none"> Self-built MongoDB MongoDB on other clouds 	Incremental synchronization supports the following modes: Replica set -> Replica set Cluster -> Cluster (the source cluster version must be 4.0 or later)	From DDS to MongoDB
	Kafka	Incremental	From DDS to Kafka
DB2 for LUW	GaussDB distributed	Full Full+Incremental	From DB2 for LUW to GaussDB Primary/Standby (To the cloud) From DB2 for LUW to GaussDB Primary/Standby (self-built - self-built)
	GaussDB primary/standby	Full Full+Incremental	From DB2 for LUW to GaussDB Distributed (To the cloud) From DB2 for LUW to GaussDB Distributed (self-built - self-built)
TiDB	GaussDB(for MySQL)	Full+Incremental	From TiDB to GaussDB(for MySQL)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
<ul style="list-style-type: none"> On-premises Microsoft SQL Server databases ECS-hosted Microsoft SQL Server databases Microsoft SQL Server-compatible databases on other clouds RDS for SQL Server 	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full+Incremental	From Microsoft SQL Server to GaussDB(DWS)
	GaussDB primary/standby	Incremental Full Full+Incremental	From Microsoft SQL Server to GaussDB Primary/Standby
	GaussDB distributed	Incremental Full Full+Incremental	From Microsoft SQL Server to GaussDB Distributed
	Microsoft SQL Server	Full+Incremental	From Microsoft SQL Server to Microsoft SQL Server
	RDS for MySQL NOTE Only whitelisted users can use this function.	Incremental Full Full+Incremental	From Microsoft SQL Server to MySQL
	GaussDB(for MySQL) NOTE Only whitelisted users can use this function.	Incremental Full Full+Incremental	From Microsoft SQL Server to GaussDB(for MySQL)
	PostgreSQL NOTE Only whitelisted users can use this function.	Incremental Full Full+Incremental	From Microsoft SQL Server to PostgreSQL
	RDS for SQL Server	Kafka NOTE Only whitelisted users can use this function.	Incremental

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
<ul style="list-style-type: none"> On-premises Microsoft SQL Server databases ECS-hosted Microsoft SQL Server databases Microsoft SQL Server-compatible databases on other clouds 	Kafka NOTE Only whitelisted users can use this function.	Incremental	From Microsoft SQL Server to Kafka (self-built - self-built)
<ul style="list-style-type: none"> On-premises MongoDB databases MongoDB databases on an ECS MongoDB database on other clouds DDS 	DDS	Full+Incremental synchronization supports the following scenarios: Replica set -> Replica set	From MongoDB to DDS
<ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds 	RDS for MariaDB	Full+Incremental	From MariaDB to MariaDB
<ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds 	RDS for MySQL	Incremental Full+Incremental	From MariaDB to MySQL
	GaussDB(for MySQL)	Incremental Full+Incremental	From MariaDB to GaussDB(for MySQL)

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
RDS for MariaDB	<ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds 	Full+Incremental	From MariaDB to MariaDB
Cassandra	GeminiDB Cassandra	Full	From Cassandra to GeminiDB Cassandra
DynamoDB on other clouds (web services)	GeminiDB Dynamo NOTE Only whitelisted users can use this function.	Full Full+Incremental	From Dynamo to GeminiDB Dynamo
<ul style="list-style-type: none"> On-premises official open-source single-node or master/standby Redis ECS-hosted official open-source single-node or master/standby Redis 	GeminiDB Redis NOTE Only whitelisted users can use this function.	Full Full+Incremental	From Redis to GeminiDB Redis
<ul style="list-style-type: none"> On-premises open-source Redis Cluster ECS-hosted open-source Redis Cluster 	GeminiDB Redis NOTE Only whitelisted users can use this function.	Full+Incremental	From Redis Cluster to GeminiDB Redis

Source DB Type	Destination DB Type	Synchronization Mode	Related Documents
GeminiDB Redis	<ul style="list-style-type: none"> On-premises open-source Codis ECS-hosted open-source Codis On-premises single-node or master/standby Redis ECS-hosted single-node or master/standby Redis GeminiDB Redis <p>NOTE Only whitelisted users can use this function.</p>	Full+Incremental	From GeminiDB Redis to Redis
	<ul style="list-style-type: none"> On-premises Redis Cluster ECS-hosted Redis Cluster <p>NOTE Only whitelisted users can use this function.</p>	Full+Incremental	From GeminiDB Redis to Redis Cluster

Other Cloud Synchronization Solutions

DRS allows you to synchronize data from other cloud databases (such as AWS RDS for MySQL and Aurora for MySQL) to Huawei Cloud databases (such as RDS for MySQL).

Table 1-2 Synchronization solutions

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
AWS	RDS for MySQL	RDS for MySQL	From MySQL to MySQL

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
AWS	RDS for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
AWS	Aurora for MySQL	RDS for MySQL	From MySQL to MySQL
AWS	Aurora for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
AWS	RDS for MariaDB	RDS for MariaDB	From MariaDB to MariaDB
AWS	RDS for PostgreSQL	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
AWS	Aurora for PostgreSQL (Non-serverless)	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
AWS	RDS for Oracle	GaussDB Primary/Standby	From Oracle to GaussDB Primary/Standby
AWS	RDS for Oracle	GaussDB Distributed	From Oracle to GaussDB Distributed
AWS	RDS for SQL Server	RDS for SQL Server	From Microsoft SQL Server to Microsoft SQL Server
AWS	RDS for SQL Server	GaussDB Primary/Standby	5.3.31-From Microsoft SQL Server to GaussDB Primary/Standby
AWS	RDS for SQL Server	GaussDB Distributed	5.3.32-From Microsoft SQL Server to GaussDB Distributed
Azure	Database for MySQL	RDS for MySQL	From MySQL to MySQL
Azure	Database for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
Azure	Database for MariaDB	RDS for MariaDB	From MariaDB to MariaDB
Azure	Database for PostgreSQL	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
Azure	SQL Server on Azure Virtual Machines	RDS for SQL Server	From Microsoft SQL Server to Microsoft SQL Server
Azure	Azure SQL Database(DB)	RDS for SQL Server	From Microsoft SQL Server to Microsoft SQL Server
Azure	Azure SQL Managed Instance(MI)	RDS for SQL Server	From Microsoft SQL Server to Microsoft SQL Server
Azure	SQL Server on Azure Virtual Machines	GaussDB Primary/Standby	5.3.31-From Microsoft SQL Server to GaussDB Primary/Standby
Azure	Azure SQL Database(DB)	GaussDB Primary/Standby	5.3.31-From Microsoft SQL Server to GaussDB Primary/Standby
Azure	Azure SQL Managed Instance(MI)	GaussDB Primary/Standby	5.3.31-From Microsoft SQL Server to GaussDB Primary/Standby
Azure	SQL Server on Azure Virtual Machines	GaussDB Distributed	5.3.32-From Microsoft SQL Server to GaussDB Distributed
Azure	Azure SQL Database(DB)	GaussDB Distributed	5.3.32-From Microsoft SQL Server to GaussDB Distributed

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
Azure	Azure SQL Managed Instance(MI)	GaussDB Distributed	5.3.32-From Microsoft SQL Server to GaussDB Distributed
Google	Cloud SQL for MySQL	RDS for MySQL	From MySQL to MySQL
Google	Cloud SQL for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Google	Cloud SQL for PostgreSQL	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
Alibaba Cloud	ApsaraDB RDS for MySQL	RDS for MySQL	From MySQL to MySQL
Alibaba Cloud	ApsaraDB RDS for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Alibaba Cloud	PolarDB for MySQL	RDS for MySQL	From MySQL to MySQL
Alibaba Cloud	PolarDB for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Alibaba Cloud	ApsaraDB RDS for MariaDB	RDS for MariaDB	From MariaDB to MariaDB
Alibaba Cloud	ApsaraDB RDS for PostgreSQL	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
Alibaba Cloud	PolarDB for PostgreSQL Standard Edition and Enterprise Edition	RDS for PostgreSQL	From PostgreSQL to PostgreSQL
Tencent Cloud	TencentDB for MySQL	RDS for MySQL	From MySQL to MySQL
Tencent Cloud	TencentDB for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Tencent Cloud	TDSQL-C for MySQL	RDS for MySQL	From MySQL to MySQL

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
Tencent Cloud	TDSQL-C for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Tencent Cloud	TencentDB for PostgreSQL	RDS for PostgreSQL	From PostgreSQL to PostgreSQL


2 Data Synchronization Topologies

DRS real-time synchronization supports multiple topology types. You can plan the topology types as required. For details, see the following content.

 **NOTE**


To ensure data consistency, do not modify the synchronization objects in the destination database.

One-to-One Real-Time Synchronization




You can create a one-to-one synchronization task.

One-to-Many Real-Time Synchronization



You need to create multiple synchronization tasks to implement one-to-many real-time synchronization. For example, to synchronize data from instance A to instances B, C, and D, you need to create three synchronization tasks.

Many-to-One Real-Time Synchronization



You need to create multiple synchronization tasks to implement many-to-one real-time synchronization. For example, to synchronize data from instances B, C, and D to instance A, you need to create three synchronization tasks.

For details about the restrictions and operation suggestions on the many-to-one scenario, see [FAQs](#).

3 To the Cloud

3.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 3-1 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)RDS for MySQL (5.5, 5.6, 5.7, 8.0)	RDS for MySQL (5.5, 5.6, 5.7, 8.0) NOTE The destination database version must be the same as or later than the source database version.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-2](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-2 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT For a full +incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.
Destination database user	The user must have the following minimum permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES The root account of the RDS for MySQL DB instance has the preceding permissions by default. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.	

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-3](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-3 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-4](#).

Table 3-4 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<<>/\'' • The source table and view names cannot contain non-ASCII characters, or the following characters: '<>/\'' • The column names in the source database tables cannot end with a backslash (\). • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • The source and destination databases cannot contain tables that have the same names but do not have primary keys. • If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • You are not advised to modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform operations (including but not limited to DDL and DML operations) on the destination database. • Do not clear binlogs on the source database. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. • Modifying MyISAM tables may cause data inconsistency. • During database name mapping, function objects, views, and stored procedures may fail to be created because they are referenced by a database table name before the mapping. DRS ignores the error reported during the creation of these objects in the full synchronization phase. As a result, these objects are lost in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • If the session variable character_set_client is set to binary, some data may include garbled characters. • During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? • You can add additional synchronization objects. • Resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down. • Some DDL statements are supported. <ul style="list-style-type: none"> – In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE,

Type	Restrictions
	<p>RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required.</p> <ul style="list-style-type: none"> - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • Do not perform the point-in-time recovery (PITR) operation on the source database. • In the incremental synchronization phase, if database name mapping is required, DDL operations (including CREATE, MODIFY, and DELETE) related to function objects, views, and stored procedures are not synchronized to the destination database. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log. • Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see Syntax Differences Between MySQL Versions. ● During table name mapping, tables on which views, stored procedures, and functions depend cannot be synchronized, and foreign key constraints of tables cannot be synchronized. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the sources and destinations are RDS instances, database mapping is required. ● During the synchronization, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database. ● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. ● Braces ({}) in a database or table name in the source database must be used in pairs. Otherwise, the table structure synchronization may fail. After data mapping, if braces ({}) in a database or name mapped to the destination database are not in pairs, the table structure may fail to be synchronized. ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● The partitioned table does not support column mapping. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. ● If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the

Type	Restrictions
	<p>JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss.</p> <ul style="list-style-type: none"> ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● After a synchronization task is created, the destination database cannot be set to read-only. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● If you create a many-to-one synchronization task, refer to Constraints and Operation Suggestions on Many-to-One Scenario. ● For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.

Prerequisites

- [You have logged in to the DRS console.](#)

- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management](#).

Procedure

This section describes how to synchronize data from a MySQL database to an RDS for MySQL database. To configure other storage engines, you can refer to the following procedures.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-1 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ 0/256

Table 3-5 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-2 Synchronization instance information

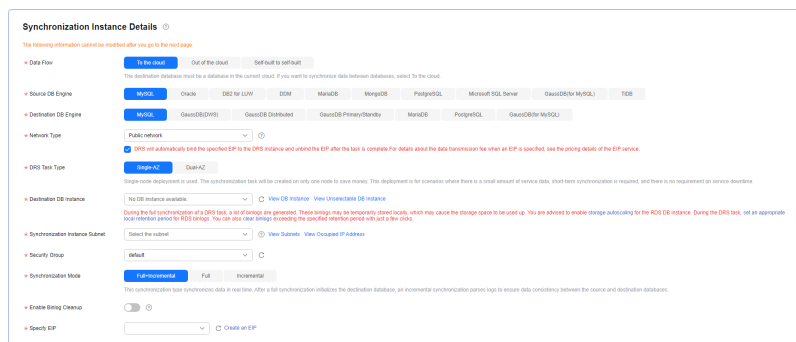


Table 3-6 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is a database in the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	<p>The RDS DB instance you created.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The destination DB instance cannot be a read replica. - The source and destination DB instances can be the same DB instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	<p>Select a security group. You can use security group rules to allow or deny access to the instance.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-3 Task type





Table 3-7 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-4 AZ</p> 

- Enterprise Project and Tags

Figure 3-5 Enterprise Project and Tags




Table 3-8 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

In different data flow scenarios, the source and destination database settings are different. Specify the required parameters based on the GUI.

- (Optional) Configuring your own DNS server

Figure 3-6 DNS Server




Table 3-9 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use the IP address of your own DNS server as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data migration.

 **NOTE**

This function is available when you need to use the IP address of your own DNS server as the source or destination database IP address.

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 3-7 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 3-10 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

- Destination database information

Figure 3-8 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-11 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 3-9 Synchronization mode

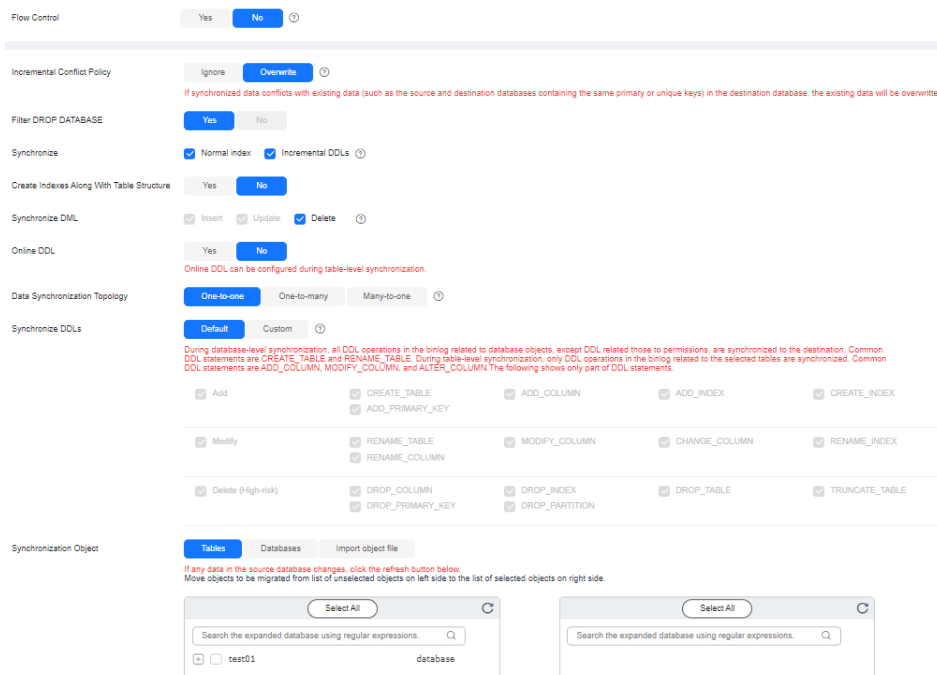
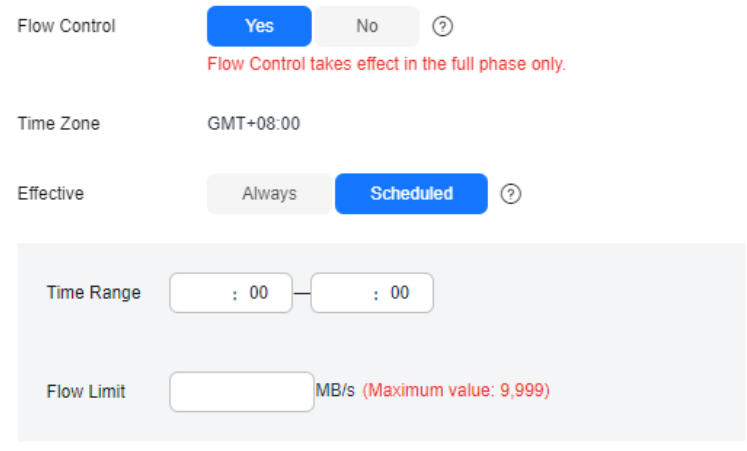
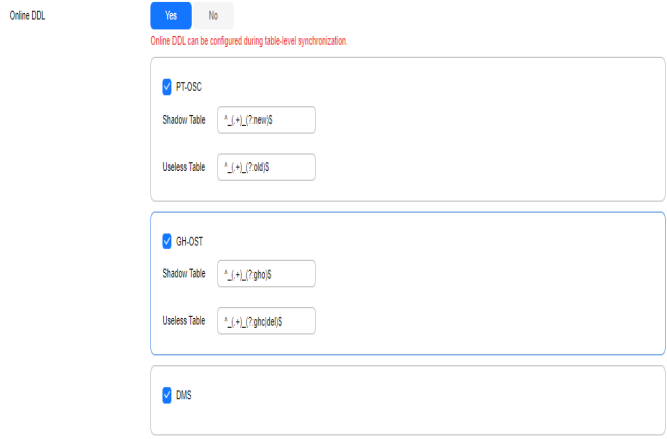



Table 3-12 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-10 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none"> • Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase. • No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OSC, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 3-11 Online DDL</p>  <ul style="list-style-type: none"> No: Table-level synchronization does not support Online DDL synchronization.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-12 Processing data



Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
gtest_order_payment_1	gtest_order_payment_1	c1	The serverName@database@table column is used	varchar(191)	Add Delete
gtest_order_payment_10	gtest_order_payment_10	c2	Default 1214	int	Add Delete
gtest_order_payment_11	gtest_order_payment_11	--	--	--	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

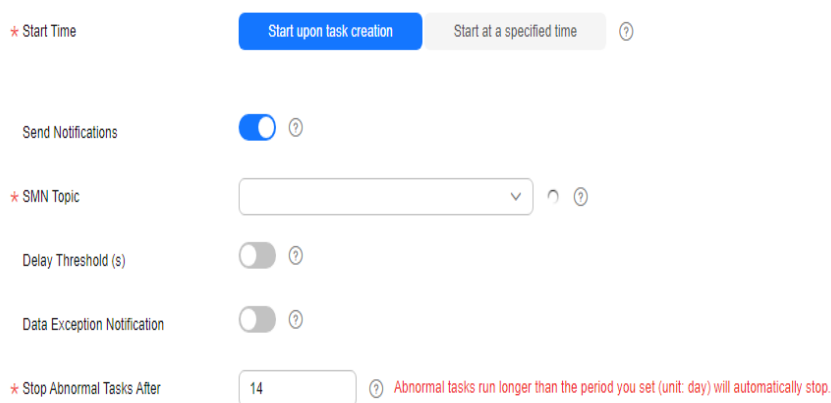
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-13 Task startup settings



* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ?

Delay Threshold (s) ?

Data Exception Notification ?


* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 3-13 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.2 From MySQL to PostgreSQL

Supported Source and Destination Databases

Table 3-14 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) RDS for MySQL (5.5, 5.6, 5.7, and 8.0) 	RDS for PostgreSQL (versions 9.5, 9.6, 10, 11, 12, 13, and 14)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-15](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Step 1

Table 3-15 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The user must have the following minimum permissions: SELECT	The user must have the following minimum permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.
Destination database user	The user must have the following minimum permissions: SELECT, INSERT, UPDATE, DELETE	The user must have the following minimum permissions: The OWNER or USERSUPER permission on the table

----End

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-16 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-16 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Only table structures, table data, and indexes can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Other database objects such as stored procedures cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-17](#).

Table 3-17 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: .<'>/\"• The column names in the source database tables cannot end with a backslash (\).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • If you change the port of the source RDS for MySQL database and the synchronization task fails, retry the task. • If the source database is not an RDS for MySQL database, the port cannot be changed. • The IP addresses, usernames, and passwords of the source and destination databases cannot be changed. • Do not clear binlogs on the source database. • Do not perform data restoration on the source database. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. • DRS writes large amount of data to the destination PostgreSQL database. As a result, the number of PostgreSQL WAL logs increases sharply, and the PostgreSQL disk space may be used up. You can disable the PostgreSQL log backup function before the full synchronization to reduce the number of WAL logs. After the synchronization is complete, enable the function. For details, see Setting an Automated Backup Policy. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> • If the session variable character_set_client is set to binary, some data may include garbled characters. • Some DDL statements are supported. <ul style="list-style-type: none"> - Table fields can be added, deleted, and modified. For example: <pre>alter table `ddl_test` add column `c2` varchar(25); alter table `ddl_test` modify column `c1` varchar(50); alter table `ddl_test` alter c1 set default 'xxx';</pre> - Table indexes can be modified. For example: <pre>alter table `ddl_test` drop primary key; alter table `ddl_test` add primary key(id); alter table `ddl_test` add index `ddl_test_uk`(id); alter table `ddl_test` drop index `ddl_test_uk`;</pre> - In table-level synchronization, you can add columns, modify columns, and add primary keys and normal indexes.

Type	Restrictions
	<ul style="list-style-type: none">- If the destination table is not found, the DDL will be ignored.- DROP_DATABASE, DROP_TABLE, TRUNCATE_TABLE, CREATE_VIEW and DROP_VIEW are not supported.- Online DDL is not supported.- During an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be synchronized on the destination database. <p>Stopping a task</p> <ul style="list-style-type: none">● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none">● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● MySQL views support syntax "as select ... from a join b where ..." but PostgreSQL does not, which may lead to the synchronization task to fail. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In PostgreSQL, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common associations: tables or views referenced by views ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● Data can be synchronized from multiple databases to one database in mapping mode. Tables with the same name cannot exist in mapped databases. ● Different types of indexes synchronized to the destination database will become B-Tree indexes. ● If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail. <ul style="list-style-type: none"> - If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs.

Type	Restrictions
	<ul style="list-style-type: none"> ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● Only data that violates the non-null constraint and data of the char or varchar type that exceeds the field length limit can be recorded. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If tables without primary keys contain large fields (BLOB, TEXT, CLOB, NCLOB, or BYTEA), data of the large fields may be inconsistent during incremental synchronization. ● If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the PostgreSQL auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-14 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ 0/256

Table 3-18 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-15 Synchronization instance details

Synchronization Instance Details

Task Name: ⓘ

Region: Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

Task Name: ⓘ

Description: ⓘ 0/256

Task Type:

Task Description:

Task Status:

Task ID:

Task Name: ⓘ

Task Description:

Task Status:

Task ID:

Table 3-19 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .

Parameter	Description
Destination DB Engine	Select PostgreSQL .
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The RDS PostgreSQL DB instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 3-16 Task type




Table 3-20 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-17 Enterprise Project and Tags




Table 3-21 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-18 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 3-22 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-19 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-23 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL instance you selected when creating the migration task and cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 3-20 Synchronization mode

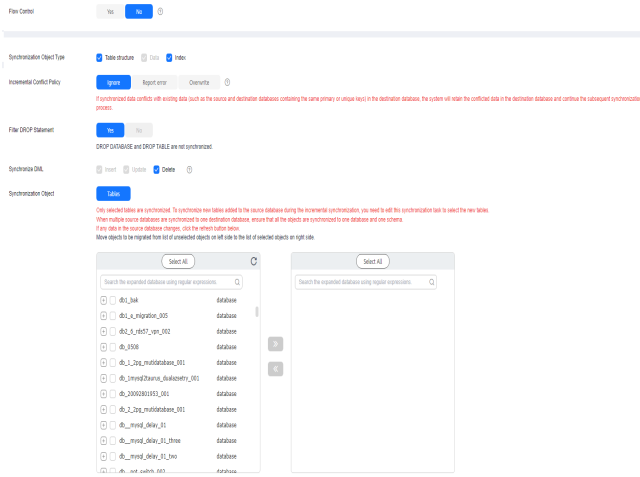
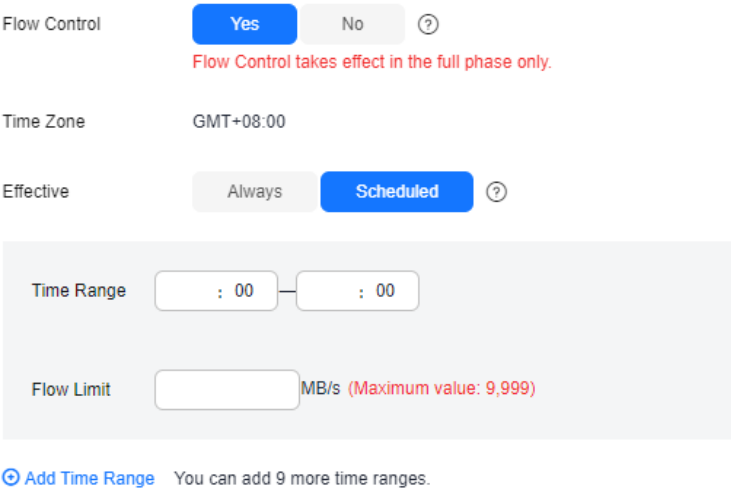



Table 3-24 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-21 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-22 Task startup settings





Table 3-25 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.3 From MySQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-26 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) RDS for MySQL (5.5, 5.6, 5.7, 8.0) 	GaussDB Distributed (MySQL-compatible) (version 1.0.0 or later)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-27](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-27 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<p>SELECT</p> <p>Statement: GRANT SELECT ON <database>.<table> to <user>;</p>	<p>SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>;</p> <p>For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-28](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-28 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• The names of the source databases, tables, and columns cannot contain non-ASCII characters or the following characters: <'>/\• The column names in the source database tables cannot end with a backslash (\).• Only table structures, table data, and indexes can be synchronized.• Only tables with primary keys can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed.• Other database objects such as stored procedures cannot be synchronized.• A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.• Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-29](#).

Table 3-29 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: .<'>/\"• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not perform data restoration on the source database. During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. During table structure synchronization of full synchronization, the auto-increment attribute of the table structure is not synchronized for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not modify or delete the usernames and permissions of the source and destination databases or change the ports of the source and destination databases. If the session variable character_set_client is set to binary, some data may include garbled characters. During incremental synchronization, the following DDL operations are supported by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, RENAME_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, ALTER_COLUMN, DROP_INDEX, and RENAME_INDEX. You can select the DDL operations to be synchronized on the object selection page as required. During an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated on the destination database. <p>Stopping a task</p>

Type	Restrictions
	<ul style="list-style-type: none">● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none">● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● By default, the MySQL-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● DRS does not proactively specify a distribution column and distribution mode, which is determined by the database itself. ● Replication tables without primary keys cannot be synchronized. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The changes to the source database cannot be synchronized to the destination database in multiple tasks at the same

Type	Restrictions
	<p>time. Otherwise, data inconsistency and synchronization failure may occur.</p> <ul style="list-style-type: none"> ● Only the database-level character set can be set for GaussDB. If the table-level and field-level character sets are set for MySQL, the table structure may fail to be created due to character length differences. ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of columns in the source database. For example, if the character set of the source database is UTF8, increase the length (byte) by three times. If the character set of the source database is UTF8MB4, increase the length (byte) by four times. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. ● If a time type is used as a primary key and the value of the time type is an invalid value of the destination database, data inconsistency or task failure may occur during incremental synchronization. ● For incremental DDL synchronization, the DDL syntax in B compatibility mode is used for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). As a result, incremental synchronization may fail. ● Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. ● If a table in the source MySQL database contains a binary field with a fixed length, the MySQL driver adds \0 to the end of the data based on the length. As a result, there may be data inconsistency after the data is synchronized to the destination GaussDB database.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000. ● If the table structure is synchronized, column names in the destination database will be converted to lowercase letters by default. If the table to be synchronized already exists in the destination database, the column names of the synchronization object table in the destination database must be in lowercase. If column names in the destination database are required to be case sensitive, choose Service Tickets > Create Service Ticket in the upper right corner of the management console to submit a service ticket. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the time type is used as a primary key and there is abnormal data (data beyond the range from 00:00:00 to 23:59:59), the value comparison result may be inconsistent.

Procedure

This section uses real-time synchronization from MySQL to GaussDB distributed as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-23 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 3-30 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-24 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** | Out of the cloud | Self-fulfill to self-fulfill

• Source DB Engine: **MySQL** | Oracle | DB2 for LUW | DDM | MariaDB | MongoDB | PostgreSQL | Microsoft SQL Server | GaussDB for MySQL | TDengine

• Destination DB Engine: MySQL | GaussDB (DBS) | **GaussDB Distributed** | GaussDB Primary/Standby | MariaDB | PostgreSQL | GaussDB for MySQL

• Network Type: Public network | Private network

• Destination DB Instance: No DB instance available | View DB Instance | View Unavailable DB Instance

• Synchronization Instance Subnet: Select the subnet | View Subnets | View Occupied IP Address

• Synchronization Mode: **Full/Incremental** | Full | Incremental

• Satisfy EP: | Create an EP

Table 3-31 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .

Parameter	Description
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB distributed instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-25 Task type




Table 3-32 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-26 Enterprise Project and Tags




Table 3-33 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-27 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 3-34 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-28 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-35 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 3-29 Synchronization Mode

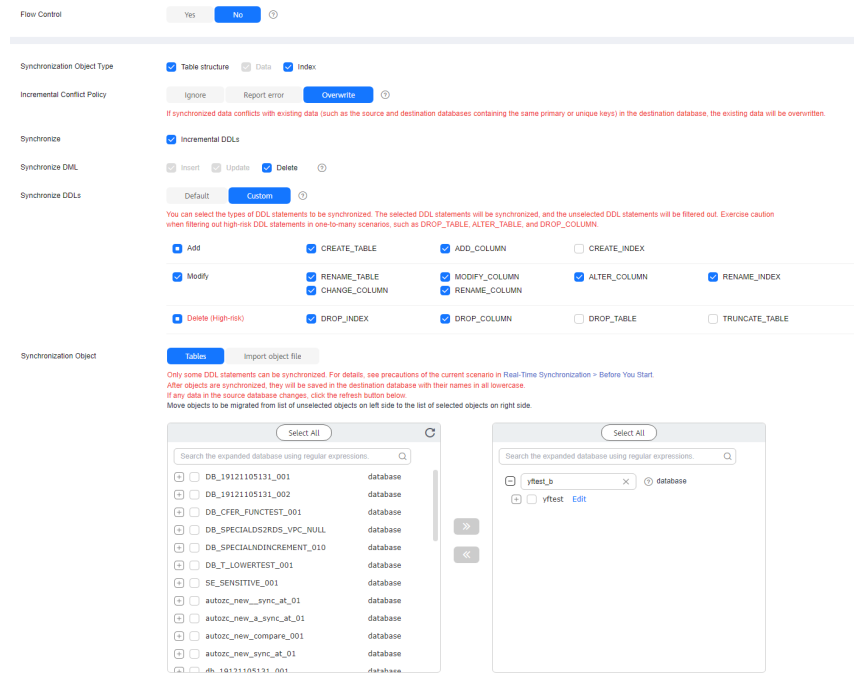
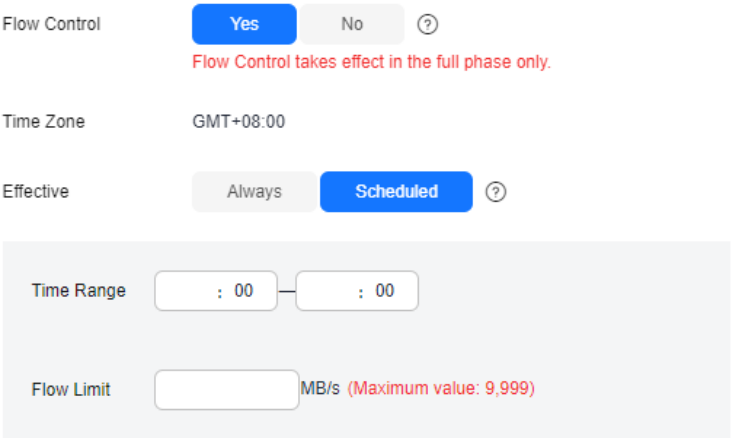



Table 3-36 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-30 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> – During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom <p>You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.</p> <p>If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-31 Processing data

Belonged Database Table	New name	Column Name	Operation Type	Type	Operation
test.mysql0	test_RS_newDatapkg_cob	-	-	-	Add
test.mysql1	test_RS_newtest1_caller	-	-	-	Add
test.mysql2	test_RS_newmysql2	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-32 Task startup settings


* Start Time: Start upon task creation Start at a specified time ?
 Send Notifications: ?
 * SMN Topic: ?
 Delay Threshold (s): ?
 Data Exception Notification: ?
 * Stop Abnormal Tasks After: ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 3-37 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.4 From MySQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-38 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • RDS for MySQL (5.5, 5.6, 5.7, 8.0) 	GaussDB Primary/Standby (B-compatible or M-compatible) (Version 1.0.0 or later)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-39](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-39 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<p>SELECT</p> <p>Statement: GRANT SELECT ON <database>.<table> to <user>;</p>	<p>SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>;</p> <p>For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> • Database-level permissions: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> • Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> • Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. • If the destination database is in M-compatible mode, the Sysadmin role must be configured. Authorization statement: ALTER USER <user> SYSADMIN; 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-40](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-40 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • The names of the source databases, tables, and columns cannot contain non-ASCII characters or the following characters: <'>/\ • The column names in the source database tables cannot end with a backslash (\). • Only table structures, table data, and indexes can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. • Other database objects such as stored procedures cannot be synchronized. • MySQL tables containing virtual columns cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-41](#).

Table 3-41 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: <'>/\" • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • The source database cannot be restored. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. • During table structure synchronization of full synchronization, the auto-increment attribute of the table structure is not synchronized for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames and permissions of the source and destination databases or change the ports of the source and destination databases. • If the session variable character_set_client is set to binary, some data may include garbled characters. • During incremental synchronization, the following DDL operations are supported by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, RENAME_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, ALTER_COLUMN, DROP_INDEX, and RENAME_INDEX. You can select the DDL operations to be synchronized on the object selection page as required. • During an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated on the destination database. <p>Stopping a task</p>

Type	Restrictions
	<ul style="list-style-type: none">● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none">● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● By default, the MySQL-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● After a table in the source database is synchronized to the destination database, the table is distributed in hash mode and cannot be replicated. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The changes to the source database cannot be synchronized to the destination database in multiple tasks at the same time. Otherwise, data inconsistency and synchronization failure may occur.

Type	Restrictions
	<ul style="list-style-type: none"> ● Only the database-level character set can be set for GaussDB. If the table-level and field-level character sets are set for MySQL, the table structure may fail to be created due to character length differences. ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of columns in the source database. For example, if the character set of the source database is UTF8, increase the length (byte) by three times. If the character set of the source database is UTF8MB4, increase the length (byte) by four times. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. ● If a time type is used as a primary key and the value of the time type is an invalid value of the destination database, data inconsistency or task failure may occur during incremental synchronization. ● For incremental DDL synchronization, the DDL syntax in B compatibility mode is used for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). As a result, incremental synchronization may fail. ● Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. ● If a table in the source MySQL database contains a binary field with a fixed length, the MySQL driver adds \0 to the end of the data based on the length. As a result, there may be data inconsistency after the data is synchronized to the destination GaussDB database. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot

Type	Restrictions
	<p>be written in strict mode, there may be data inconsistency during synchronization.</p> <ul style="list-style-type: none"> ● If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000. ● If the table structure is synchronized, column names in the destination database will be converted to lowercase letters by default. If the table to be synchronized already exists in the destination database, the column names of the synchronization object table in the destination database must be in lowercase. If column names in the destination database are required to be case sensitive, choose Service Tickets > Create Service Ticket in the upper right corner of the management console to submit a service ticket. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the time type is used as a primary key and there is abnormal data (data beyond the range from 00:00:00 to 23:59:59), the value comparison result may be inconsistent.

Procedure

This section uses real-time synchronization from MySQL to GaussDB primary/standby as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-33 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Region dropdown] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Project dropdown]

* Task Name: DRS-5678

Description: [Description text area] 0/256

Table 3-42 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-34 Synchronization instance information

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** (Out of the cloud, Default to self built)

• Source DB Engine: **MySQL** (Oracle, DB2 for LUW, DDM, MariaDB, MaxGDB, PostgreSQL, Microsoft SQL Server, GaussDB for MySQL, TDB)

• Destination DB Engine: **GaussDB Primary/Standby** (GaussDB Standby, GaussDB MySQL)

• Network Type: **Public network**

• Destination DB Instance: [DB Instance available] [View DB Instance](#) [View Unavailable DB Instance](#)

• Synchronization Instance Subnet: [Select the subnet] [View Subnets](#) [View Occupied IP Address](#)

• Synchronization Mode: **Full** (Full, Incremental)

• Specify EP: [Specify EP] [Create an EP](#)

Table 3-43 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB Primary/Standby .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB primary/standby instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-35 Task type




Table 3-44 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-36 Enterprise Project and Tags




Table 3-45 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-37 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 3-46 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-38 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-47 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 3-39 Synchronization Mode

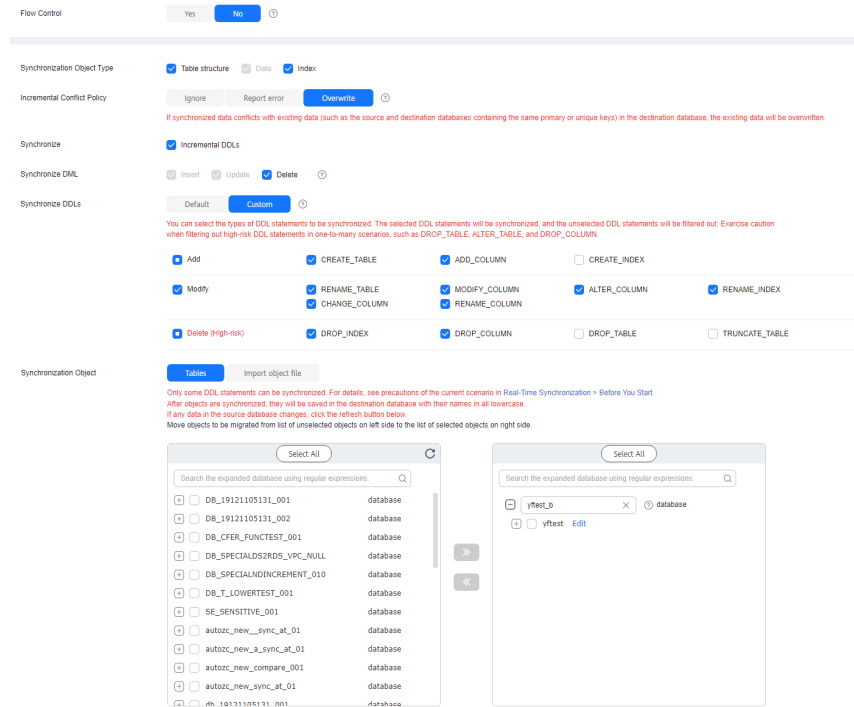
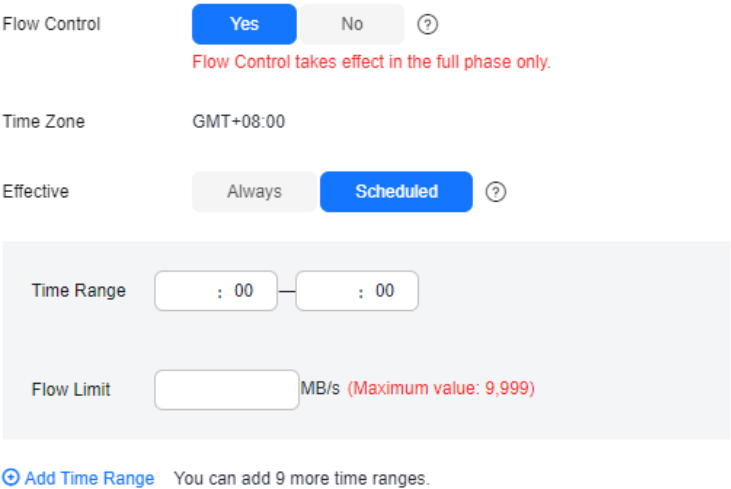



Table 3-48 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-40 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten.
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> – During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom <p>You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.</p> <p>If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-41 Processing data

Belonged Database Table	New Name	Column Name	Operation Type	Type	Operation
gboot_order_payment_1	gboot_order_payment_1	c1	The serverName@database@table column is used	varchar(19)	Add Detail
gboot_order_payment_10	gboot_order_payment_10	c2	Default 1234	int	Add Detail
gboot_order_payment_11	gboot_order_payment_11	--	--	--	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-42 Task startup settings

Table 3-49 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.5 From MySQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-50 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) • RDS for MySQL (5.5, 5.6, 5.7, 8.0) 	GaussDB(DWS) (versions 8.1.3 and 8.2.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-51](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-51 Database account permission

Type	Full, Incremental, and Full+Incremental
Source database user	The user must have the following minimum permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT
Destination database user	Each table in the destination database must have the INSERT, SELECT, UPDATE, DELETE, CONNECT and CREATE permissions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-52](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-52 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, common indexes (B-Tree indexes), and constraints (primary key, empty, and not-null) can be synchronized. • MyISAM and InnoDB tables can be synchronized. • Comment is supported in full synchronization mode, but not supported in incremental synchronization mode. • Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, and unique indexes cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. • Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-53](#).

Table 3-53 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: .<'>/"• The column names in the source database tables cannot end with a backslash (\).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • If you change the port of the RDS source database and the synchronization task fails, retry the task. • During real-time synchronization, if the source is not RDS, the port cannot be changed. • During real-time synchronization, the IP address, account, and password cannot be changed. • Do not perform data restoration on the source database. • Binlogs cannot be forcibly cleared. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • If the session variable character_set_client is set to binary, some data may include garbled characters. • During an incremental synchronization, tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated on the destination database. • Tables without primary keys created during database-level synchronization must contain one of the following. For details, see GaussDB(DWS) Parameter Description. Integer types: TINYINT, SMALLINT, INT, BIGINT, and NUMERIC/DECIMAL Character types: CHAR, BPCHAR, VARCHAR, VARCHAR2, and NVARCHAR2 Date/time types: DATE, TIME, TIMETZ, TIMESTAMP, TIMESTAMPTZ, INTERVAL, and SMALLDATETIME <p>For a table without a primary key, a column that can be used as a distribution column is specified as the distribution column. If all column types cannot be used as distribution columns, the table fails to be created and DRS synchronization is interrupted.</p> <ul style="list-style-type: none"> • During incremental synchronization, DDL operations (such as ALTER TABLE ddl_test ADD COLUMN c2 AFTER/FIRST c1) for adding columns to a specified position are not supported. DRS will delete the AFTER/FIRST attribute, which may cause column sequence inconsistency.

Type	Restrictions
	<ul style="list-style-type: none"> • During incremental synchronization, executing non-idempotent DDL statements (for example, ALTER TABLE ddl_test ADD COLUMN c3 timestamp default now();) may cause data inconsistency. • During an incremental synchronization, database-level synchronization does not support online DDL, and table-level synchronization supports only online DDL generated by Alibaba Cloud DMS. • During incremental synchronization, the following DDL operations are supported by default: <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, ALTER_COLUMN, DROP_CONSTRAINT, ADD_CONSTRAINT, CREATE_INDEX, DROP_INDEX, RENAME_INDEX, and RENAME_COLUMN. You can select the DDL operations to be synchronized on the object selection page as required. - If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. - In many-to-one scenarios, you are advised to synchronize ADD_COLUMN only. Other DDL synchronization may cause task failures or data inconsistency due to destination table changes. - In many-to-one scenarios, when running ADD_COLUMN, ensure that the types of columns added to each table are the same. Otherwise, the task may fail. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - When an index is creating using a DDL statement, if the table name and index name in the statement contain more than 63 characters, duplicate names may occur. As a result, the index fails to be created. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. - When a DDL operation is performed in the incremental phase, if the destination table is not found, the DDL operation will be ignored. - In the incremental phase, if CHANGE COLUMN is performed in the source database to modify a column and the column is a distribution column in destination GaussDB(DWS), the statement will be ignored because GaussDB(DWS) does not support distribution column modification.

Type	Restrictions
	<ul style="list-style-type: none"> - In the incremental phase, the RENAME INDEX operation is not supported because the index rules of the source and destination databases are different. - In the incremental phase, you are not advised to combine CHAR(0) with other characters for data synchronization. For example, inserting CONCAT('a',CHAR(0),'b') may cause data inconsistencies. - In the incremental phase, you are not advised to combine CHAR(34) (double quotation marks) and CHAR(92) (backslash) with other characters for JSON data synchronization. For example, inserting JSON_OBJECT('\$\$.0', CONCAT('a', CHAR(34), 'b')) may cause data inconsistencies because an extra backslash (\) will be added after it is synchronized to the destination end. <p>Stopping a task</p> <ul style="list-style-type: none"> ● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log. ● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● Associated objects must be synchronized at the same time to avoid synchronization failure due to missing associated objects. Common association: Indexes reference tables. ● The source database cannot be restored. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● If there is a unique key when a table without a primary key is synchronized, data may fail to be written due to data conflicts. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If tables without primary keys contain large fields (BLOB, TEXT, CLOB, NCLOB, or BYTEA), data of the large fields may be inconsistent during incremental synchronization. ● If the source is an RDS for MySQL instance and the destination is a GaussDB(DWS) instance, multiple source tables can be mapped to one destination table. For details, see From MySQL to GaussDB(DWS). ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● MySQL 5.7 filters CHECK constraints. If the source database version is 5.7, you are not advised to use CHECK constraints, which may cause data inconsistencies. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● After being synchronized to the destination GaussDB(DWS) database, the unique key of the MySQL database is changed to a common constraint. The primary key remains unchanged. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique

Type	Restrictions
	<p>constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail.</p> <ul style="list-style-type: none"> ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. During incremental partition table synchronization, partition tables are created based on the types supported by GaussDB(DWS), which may fail. ● When the timestamp data is to be synchronized in full synchronization, the on update current_timestamp syntax in the default value will not be synchronized to the destination GaussDB(DWS). ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> – If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out.

Type	Restrictions
	<ul style="list-style-type: none"> - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. • The SQL statements in the DDL for incremental synchronization come from the native SQL statements executed by users. The table structure syntax for full synchronization is the SQL statements processed by the SQL syntax layer. The types or constraints of these SQL statements are different. As a result, the table structures or field types created in full synchronization and incremental synchronization may be different. For details, see Suggestions on Synchronizing Data to GaussDB(DWS). • If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000.

Procedure

This section uses many-to-one synchronization from RDS for MySQL to GaussDB(DWS) as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-43 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.

The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

▼

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

▼

* Task Name

DRS-5678
⊙

Description

⊙

0/256

Table 3-54 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-44 Synchronization instance information

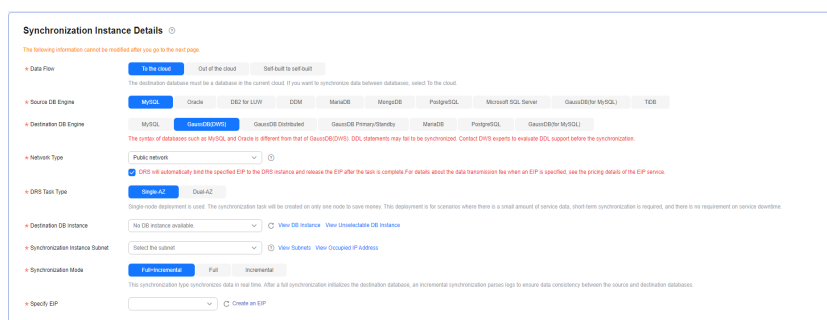


Table 3-55 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB(DWS) .

Parameter	Description
Network Type	<p>The VPC network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	An available GaussDB(DWS) instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-45 Task type



Table 3-56 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-46 AZ</p>

- Enterprise Project and Tags

Figure 3-47 Enterprise Project and Tags




Table 3-57 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-48 Source database information

Source Database

DRS migrates only some key parameters to the destination database. For the other parameters that cannot be migrated, you need to use parameter templates to configure them on the destination database.

Database Type: self-built database RDS DB instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

SSL Connection:

Table 3-58 Source database settings

Parameter	Description
Database Type	Select Self-built database or RDS DB instance .
DB Instance Name	Select the RDS DB instance to be synchronized.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-49 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-59 Destination database settings

Parameter	Description
DB Instance Name	The default value is the GaussDB(DWS) instance selected for creating the migration task. It cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization object. Click **Next**.

Figure 3-50 Synchronization Mode

Flow Control: Yes **No** ⓘ

Synchronization Object Type: Table structure Data Index
When you manually create a table structure in the destination database, for details about the data type, see Mapping Data Types.

Incremental Conflict Policy: Ignore Report error **Overwrite** ⓘ
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Abnormal Data Policy: **Continue** Failed ⓘ
A small amount of abnormal data will not cause the synchronization of all tables to be stopped. If there is abnormal data, dirty data is automatically recorded. You can view the dirty data in the abnormal data module.

Online DDL: Yes **No** ⓘ

Synchronize DML: Insert Update Delete ⓘ

Data Synchronization Topology: **One-to-one** One-to-many Many-to-one ⓘ

Synchronize DDLs: **Default** **Custom** ⓘ
You can select the types of DDL statements to be synchronized. The selected DDL statements will be synchronized, and the unselected DDL statements will be filtered out. Exercise caution when filtering out high-risk DDL statements in one-to-many scenarios, such as DROP_TABLE, ALTER_TABLE, and DROP_COLUMN.

<input checked="" type="checkbox"/> Add	<input checked="" type="checkbox"/> CREATE_TABLE	<input checked="" type="checkbox"/> ADD_COLUMN	<input checked="" type="checkbox"/> ADD_CONSTRAINT	<input checked="" type="checkbox"/> CREATE_INDEX
<input checked="" type="checkbox"/> Modify	<input checked="" type="checkbox"/> RENAME_TABLE	<input checked="" type="checkbox"/> MODIFY_COLUMN	<input checked="" type="checkbox"/> ALTER_COLUMN	<input checked="" type="checkbox"/> RENAME_INDEX
<input checked="" type="checkbox"/> Delete (High-risk)	<input checked="" type="checkbox"/> DROP_CONSTRAINT	<input checked="" type="checkbox"/> DROP_INDEX	<input type="checkbox"/> DROP_COLUMN	<input type="checkbox"/> DROP_TABLE
	<input type="checkbox"/> TRUNCATE_TABLE	<input type="checkbox"/> DROP_PARTITION		

Synchronization Object: **Tables** Databases Import object file
Only some DDL statements can be synchronized. For details, see precautions of the current scenario in Real-Time Synchronization > Before You Start.
 If any data in the source database changes, click the refresh button below.
 Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

Search the expanded database using regular expressions

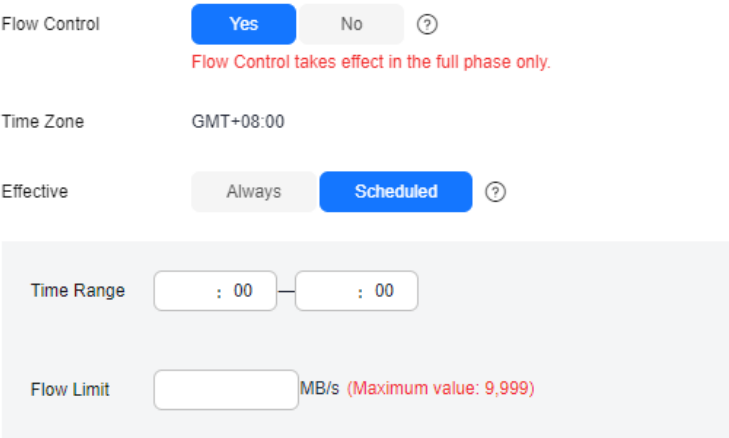
- cubc_trade_00 database
- db01 database
- db01qt database
- db01pyj database
- db01pyj1 database
- db01pyj11 database
- db01pyj1111 database
- db01pyj1111dd database
- db01pyj112 database
- db01pyj112d4 database
- db01pyj112zh database

Select All

Search the expanded database using regular expressions


- auto_05m01_db_001 database
- db001 Edit

Table 3-60 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-51 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. <p>The table structure of a GaussDB(DWS) database has service logic. You are advised to create a table structure under the guidance of GaussDB(DWS) experts based on service requirements and ensure that the table names, column names, and column types are correct. Otherwise, data synchronization may fail and data synchronization precision may be lost.</p>
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Abnormal Data Policy	<p>Differences between the source and destination databases may cause some data to fail to be written. In this case, there may be abnormal data. Select this option based on how well abnormal data can be tolerated.</p> <ul style="list-style-type: none"> • Continue: A small amount of abnormal data will not cause the synchronization of all data to be stopped. If there is abnormal data, dirty data is automatically recorded. You can view the dirty data in the abnormal data module. • Failed: If there is an exception during a task, the task fails and does not continue.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Data Synchronization Topology	<p>Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>
Synchronize DDLs	<p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. When selecting synchronization objects, you can also map multiple tables to one table. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> - When you map multiple tables to one table, use the additional column to process data to avoid data conflicts. - In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. - The structure of the source database tables to be mapped must be the same as that of the destination database table. • For details about how to import an object file, see Importing Synchronization Objects. If you select Import object file, you can map multiple tables to one table and import up to 10,000 tables. If the table name or the mapping rule is too long, the number of tables to be imported will be affected. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • Object names will be converted to lowercase letters after being synchronized to the destination database. Therefore, the selected source database tables cannot contain tables with the same name but different letter cases. Otherwise, the synchronization fails. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. • In incremental synchronization at the database level, do not create tables with the same name but different letter cases in the source database. Otherwise, one of the tables cannot be synchronized.

Step 5 On the **Processing Data** page, filter data or add additional columns for the table object to be processed, and click **Next**.

- If you need to set data filtering, click **Data Filtering** and set related filtering rules.
- If you need to add additional columns, click the **Additional Columns** tab, click **Add** in the **Operation** column, and enter the column name and the operation type.

For details about related operations, see [Processing Data](#).

Figure 3-52 Processing data

Data Filtering Additional Columns

You can use additional columns to avoid data conflicts in many-to-one operations.

Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
db001_mysql_01	db001_mysql_01	-	-	-	Add
db001_mysql_02	db001_mysql_02	-	-	-	Add
db001_mysql_03	db001_mysql_03	-	-	-	Add
db001_mysql_04	db001_mysql_04	-	-	-	Add
db001_mysql_datatype	db001_mysql_datatype	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 3-53 Pre-check

Check success rate 100% All checks must pass before you can continue. If any check requires confirmation, check and confirm the results before proceeding to the next step.

Check Item	Check Result
Database parameters	
Whether the destination database users (schemas) and tables exist.	Passed
Whether the source and destination database character sets are consistent	Passed
Whether the source database name is valid	Passed
Whether the source database table contains unsupported data types	Passed
Whether the source database contains replication tables	Passed
Whether the source database contains compression tables	Passed
Whether the source database contains column tables	Passed
Whether the source database schema name is valid	Passed
Whether the source database table name is valid	Passed

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-54 Task startup settings





Table 3-61 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.6 From MySQL to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-62 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) RDS for MySQL (5.5, 5.6, 5.7, 8.0) 	GaussDB(for MySQL) Primary/Standby (version 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-63](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-63 Database account permission

Type	Full+Incremental Synchronization and Incremental Synchronization
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.
Destination database user	The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-64](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-64 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. • Events and triggers cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-65](#).

Table 3-65 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The names of the source databases, tables, and views cannot contain non-ASCII characters, or the following characters: . > < / \ ' " • The column names in the source database tables cannot end with a backslash (\). • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. Modifying MyISAM tables may cause data inconsistency. During database name mapping, function objects, views, and stored procedures may fail to be created because they are referenced by a database table name before the mapping. DRS ignores the error reported during the creation of these objects in the full synchronization phase. As a result, these objects are lost in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> If the session variable character_set_client is set to binary, some data may include garbled characters. If database name mapping is required, DDL operations (including CREATE, MODIFY, and DELETE) related to function objects, views, and stored procedures are not synchronized to the destination database. Some DDL operations are supported. <ul style="list-style-type: none"> In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The source and destination databases cannot contain tables that have the same names but do not have primary keys. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique

Type	Restrictions
	<p>constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail.</p> <ul style="list-style-type: none"> ● The source database does not support point-in-time recovery (PITR). ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● After a task is created, the destination database cannot be set to read-only. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. ● If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail.

Procedure

This section uses real-time synchronization from MySQL to GaussDB(for MySQL) as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-55 Synchronization task information

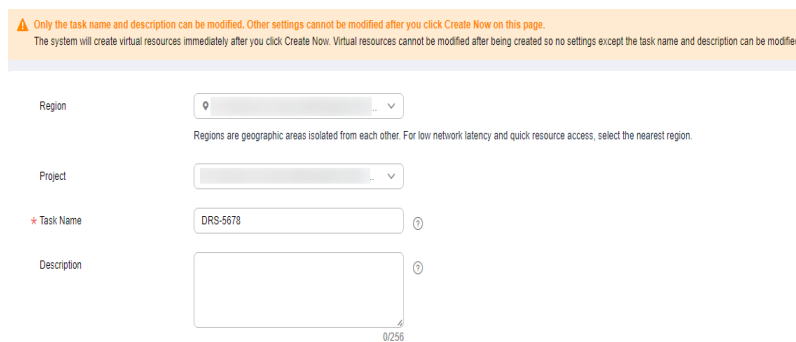


Table 3-66 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-56 Synchronization instance information

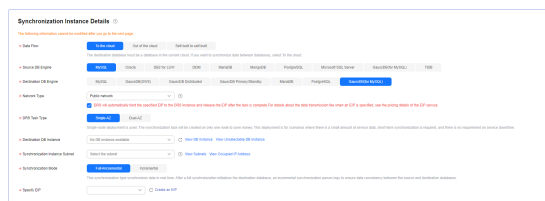


Table 3-67 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB(for MySQL) .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	An available GaussDB(for MySQL) instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental <p>This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.</p> <ul style="list-style-type: none"> – Incremental <p>Through log parsing, incremental data generated on the source database is synchronized to the destination database.</p> <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-57 Task type





Table 3-68 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-58 AZ</p> 

- Enterprise Project and Tags

Figure 3-59 Enterprise Project and Tags




Table 3-69 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-60 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 3-70 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-61 Destination database information

Destination Database

DB Instance Name

Database Username


Database Password 

Table 3-71 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 3-62 Synchronization Mode

Flow Control Yes No ⓘ

Incremental Conflict Policy Ignore Overwrite ⓘ
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Filter DROP DATABASE Yes No

Synchronize Normal index ⓘ

Create Indexes Along With Table Structure Yes No

Synchronize DML Insert Update Delete ⓘ

Online DDL Yes No
Online DDL can be configured during table-level synchronization.

Data Synchronization Topology One-to-one One-to-many Many-to-one ⓘ

Synchronize DDLs Default Custom ⓘ
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. The following shows only part of DDL statements.

<input checked="" type="checkbox"/> Add	<input checked="" type="checkbox"/> CREATE_TABLE <input checked="" type="checkbox"/> ADD_PRIMARY_KEY	<input checked="" type="checkbox"/> ADD_COLUMN	<input checked="" type="checkbox"/> ADD_INDEX	<input checked="" type="checkbox"/> CREATE_INDEX
<input checked="" type="checkbox"/> Modify	<input checked="" type="checkbox"/> RENAME_TABLE <input checked="" type="checkbox"/> RENAME_COLUMN	<input checked="" type="checkbox"/> MODIFY_COLUMN	<input checked="" type="checkbox"/> CHANGE_COLUMN	<input checked="" type="checkbox"/> RENAME_INDEX
<input checked="" type="checkbox"/> Delete (High-risk)	<input checked="" type="checkbox"/> DROP_COLUMN <input checked="" type="checkbox"/> DROP_PRIMARY_KEY	<input checked="" type="checkbox"/> DROP_INDEX <input checked="" type="checkbox"/> DROP_PARTITION	<input checked="" type="checkbox"/> DROP_TABLE	<input checked="" type="checkbox"/> TRUNCATE_TABLE

Synchronization Object Tables Databases Import object file

Only selected tables are synchronized. To synchronize new tables added to the source database during the incremental synchronization, you need to edit this synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

ⓘ

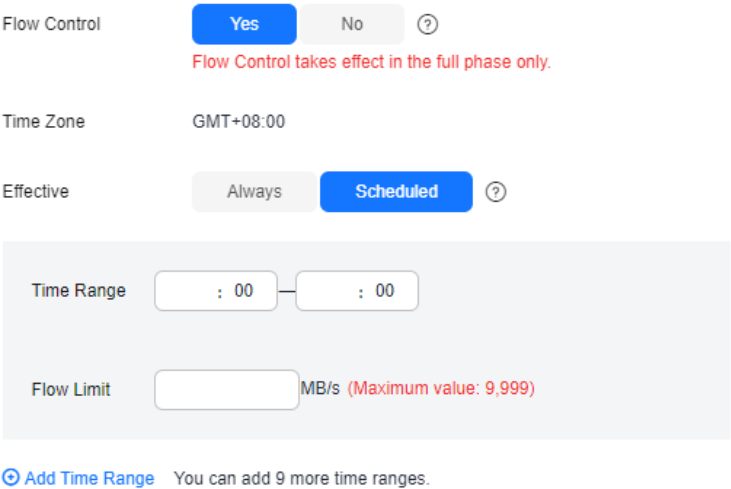
Search the expanded database using regular expressions.

test01 database

ⓘ

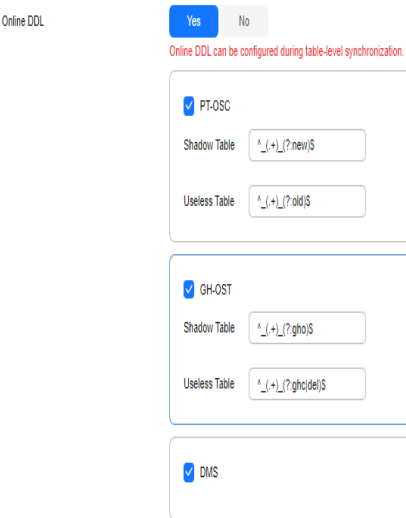
Search the expanded database using regular expressions.


Table 3-72 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-63 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Whether to synchronize normal indexes.</p> <p>By default, DRS synchronizes the primary key or unique index. A normal index refers to an index other than the primary key or unique index. If you select normal index, all indexes will be synchronized. If you do not select normal index, only the primary key and unique index will be synchronized.</p>
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none"> ● Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase. ● No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>
Synchronize DDLs	<p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OSC, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 3-64 Online DDL</p>  <p>The screenshot shows the 'Online DDL' configuration interface. At the top, there is a toggle switch for 'Online DDL' with 'Yes' selected. Below the toggle, a red text note states 'Online DDL can be configured during table-level synchronization.' There are three main configuration sections: <ul style="list-style-type: none"> PT-OSC: Includes a checked checkbox, a 'Shadow Table' input field with the default value '^_(+)_?.new/\$', and a 'Useless Table' input field with the default value '^_(+)_?.old/\$'. GH-OST: Includes a checked checkbox, a 'Shadow Table' input field with the default value '^_(+)_?.gho/\$', and a 'Useless Table' input field with the default value '^_(+)_?.gho(del)/\$'. DMS: Includes a checked checkbox. </p> <ul style="list-style-type: none"> No: Table-level synchronization does not support Online DDL synchronization.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-65 Task startup settings

Table 3-73 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.7 From MySQL to MariaDB

Supported Source and Destination Databases

Table 3-74 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0) RDS for MySQL (5.5, 5.6, 5.7, 8.0) 	RDS for MariaDB 10.5

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-75](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-75 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.
Destination database user	The user must have the following minimum permissions: The root account of RDS for MariaDB has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-76](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-76 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • The table structure of the synchronization object in the source database cannot contain the table identifier ENCRYPTION. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-77](#).

Table 3-77 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> ● The source database names cannot contain non-ASCII characters, or the following characters: .></\''" ● The column names in the source database tables cannot end with a backslash (\). ● The table names of the source database cannot contain non-ASCII characters, or the following special characters: ></\ ● The source database names cannot contain non-ASCII characters, or the following characters: . > < / \ ' `" ● The table names of the source database cannot contain non-ASCII characters, or the following special characters: > < / \ ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● If the source MySQL database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • A real-time synchronization task may fail due to the change of the username and password of the source or destination database. If it happens, rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization. • To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. • Modifying MyISAM tables may cause data inconsistency. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • If the session variable character_set_client is set to binary, some data may include garbled characters. • Some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out.

Type	Restrictions
	<ul style="list-style-type: none">- If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out.- You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.• During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization?• You can add additional synchronization objects. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the enforce_storage_engine parameter of the destination MariaDB database is set to InnoDB, DRS cannot synchronize the table structure and data whose storage engine is MyISAM to the destination MariaDB database. To synchronize table data whose storage engine is MyISAM, create a table structure on the destination database. (The storage engine can only be set to InnoDB due to the value restriction of the enforce_storage_engine parameter.) ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions? ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The source database does not support point-in-time recovery (PITR).

Type	Restrictions
	<ul style="list-style-type: none"> ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The partitioned table does not support column mapping. ● After a task is created, the destination database cannot be set to read-only. ● During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. ● If the source database version is MySQL 8.0, some collation character sets (such as utf8mb4_0900_as_ci, utf8mb4_0900_as_cs, utf8mb4_0900_bin and utf8mb4_cs_0900_ai_ci that support Unicode 9.0) are not supported by the destination database. ● Due to engine and version differences, the following MySQL functions are not supported in MariaDB: MBRCOVEREDBY, ST_BUFFER_STRATEGY, ST_GeoHash, ST_IsValid, ST_LatFromGeoHash, ST_LongFromGeoHash, ST_PointFromGeoHash, ST_SIMPLIFY, ST_VALIDATE, (8.0)JSON_ARRAYAGG, JSON_OBJECTAGG, JSON_PRETTY, JSON_STORAGE_FREE, JSON_STORAGE_SIZE and JSON_TABLE. ● The destination MariaDB database does not support VALIDATION during partition exchange. Do not use the verification syntax when executing partition exchange SQL statements in the source MySQL database. For example: ALTER TABLE t1 EXCHANGE PARTITION p0 WITH TABLE t2 WITH VALIDATION; ● Do not write characters such as \n, \t, and \r to the JSON data of the source MySQL database during incremental synchronization. Otherwise, there may be data inconsistencies. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

This section describes how to synchronize data from a MySQL database to an RDS for MariaDB database. To configure other storage engines, you can refer to the following procedures.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-66 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 3-78 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-67 Synchronization instance information

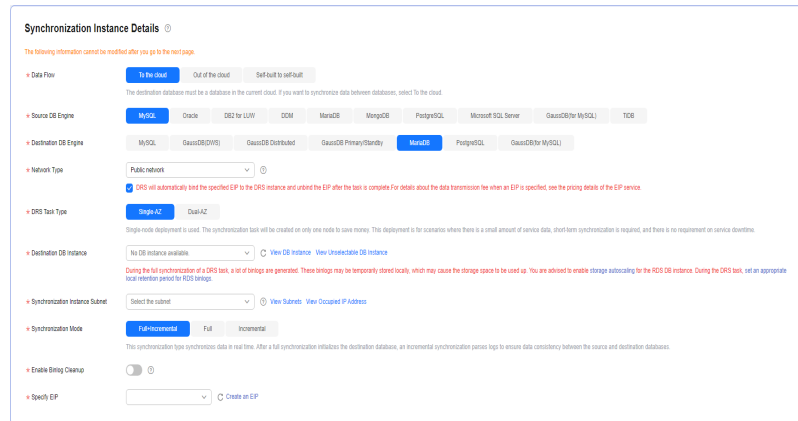


Table 3-79 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is a database in the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MariaDB .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	<p>The RDS for MariaDB instance you created.</p>
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-68 AZ

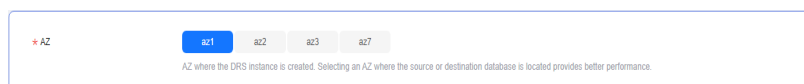



Table 3-80 Task AZ

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-69 AZ</p> 

- Enterprise Project and Tags

Figure 3-70 Enterprise Project and Tags




Table 3-81 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 3-71 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 3-82 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.

Parameter	Description
Database Password	The password for the database username.
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

- Destination database information

Figure 3-72 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-83 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MariaDB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 3-73 Synchronization Mode

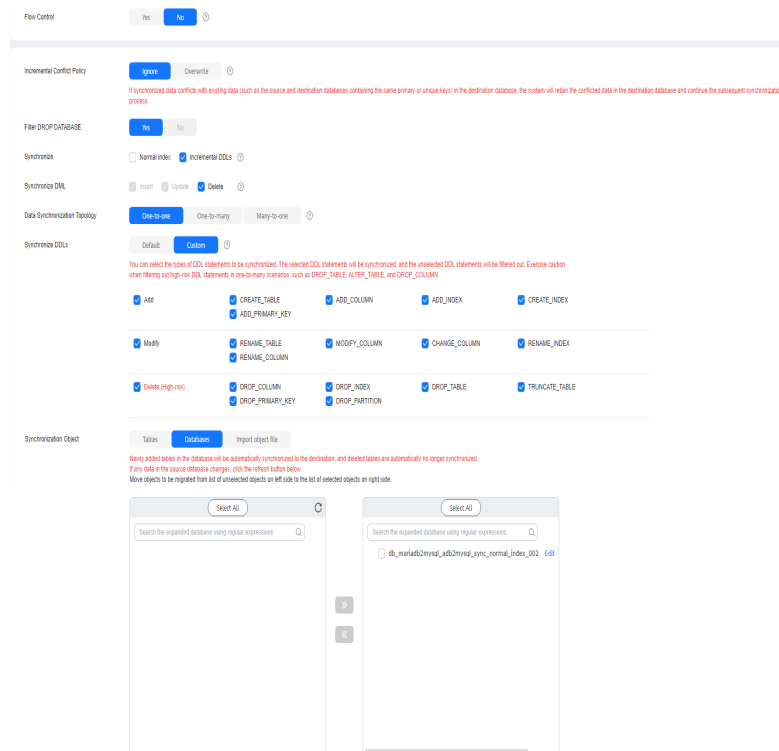
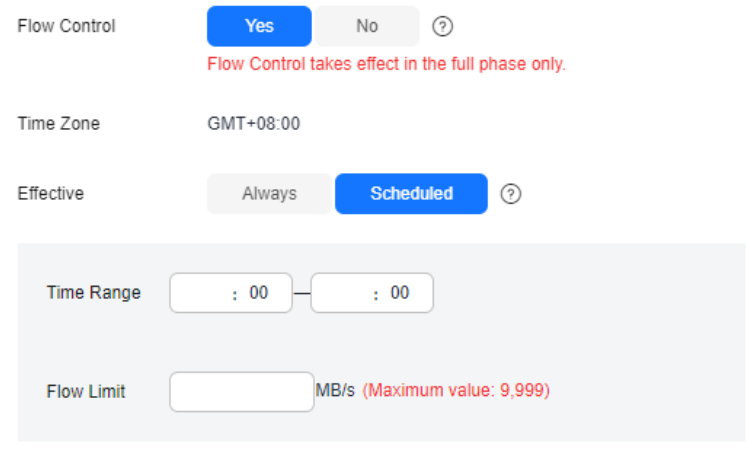



Table 3-84 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-74 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>

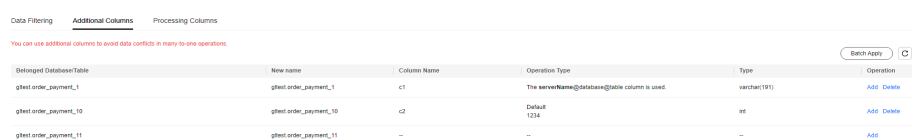
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-75 Processing data



Belonged Database/Table	New Name	Column Name	Operation Type	Type	Operation
gtest.order_payment_1	gtest.order_payment_1	c1	The serverName@database@table column is used.	varchar(191)	Add Delete
gtest.order_payment_10	gtest.order_payment_10	c10	Default 1234	int	Add Delete
gtest.order_payment_11	gtest.order_payment_11	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-76 Task startup settings

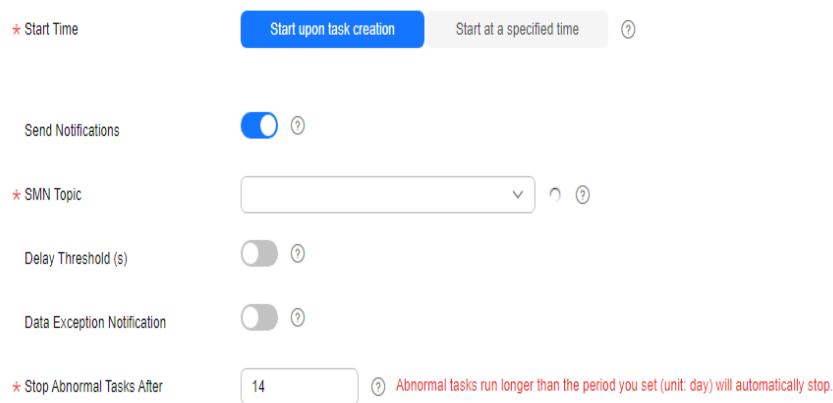



Table 3-85 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.8 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 3-86 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises databases (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13, 14 and 15) ECS-hosted databases (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13, 14 and 15) Other cloud databases (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13, 14 and 15) RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14 and 15) 	RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14 and 15) NOTE The major version of the destination database must be the same as or later than that of the source database.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-87](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-87 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, and the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords)	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords), the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> 1. Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the pg_hba.conf file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> ● Database-level: <ul style="list-style-type: none"> - If the destination database is not PostgreSQL, the CREATEDB permission is required. - If the destination database is PostgreSQL, the CONNECT and CREATE permissions on PostgreSQL databases and the USAGE and CREATE permissions on public schemas are required. ● Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the object are required. ● Synchronization user: The CREATEROLE permission is required. ● Synchronization user permissions: The default privilege cannot be modified. Otherwise, the object permissions of the destination database may be inconsistent with those of the source database. <p>NOTE</p> <ul style="list-style-type: none"> ● To synchronize event triggers, text search parsers, and text search templates, the destination database version must be RDS for PostgreSQL 11.11 or later, and the destination database user must be user root or a member of user root. ● If the destination database version is PostgreSQL 15 and no database is created on the destination database for database-level synchronization and table-level synchronization, full or incremental synchronization may fail due to lack of the USAGE and CREATE permissions of the public schema. 	

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-88 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-88 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● Scope of full synchronization <ul style="list-style-type: none"> - The following objects are supported in the database-level synchronization: schemas, tables, indexes, constraints, views, materialized views, sequences, stored procedures, rules, triggers, foreign keys, sorting rules, plug-ins, code conversion information, aggregate functions, operators, statistics extension, conversion information, text search configurations, functions, data types, type conversion, users, event triggers, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, users, and common indexes can be synchronized. During object file import, tables can be synchronized. - Objects that are not supported: system schemas (schemas starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random), system catalogs, system users, tablespaces, foreign-data wrappers, foreign servers, user mappings, publications, and subscriptions ● Instance-level synchronization is not supported.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-89](#).

Table 3-89 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> ● The partition table trigger of the source database cannot be set to disable. ● To perform incremental synchronization: The pg_hba.conf file of the source database contains the following configuration: host replication all 0.0.0.0/0 md5 ● The test_decoding plug-in must be installed on the source database in advance. ● Triggers with the same name cannot exist in the source database.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <ul style="list-style-type: none"> Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, data may be inconsistent or the task may fail. During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail. Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. Some DDL statements can be synchronized, including TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA, CREATE TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW. During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, COMMENT ON

Type	Restrictions
	<p>COLUMN, COMMENT ON TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p> <ul style="list-style-type: none"> • Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE</p> <ul style="list-style-type: none"> • The source database captures DDL statements using event triggers and records them in specific tables, so you need to create event triggers and functions in the source database in advance. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL. • Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.) <p>Synchronization comparison</p> <ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. <ul style="list-style-type: none"> - The destination database sequence value is automatically reset. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. If users are synchronized, the user memberships are automatically synchronized after the task is complete. - When a full synchronization task is complete, the sequence values are compared based on logical consistency. That is, if an auto-increment sequence value is used only when the sequence value of the destination database is greater than or equal to that of the source database, or an auto-decrement sequence value is used only when the sequence value of the destination database is less than or equal to that of the source database, the sequence values in the

Type	Restrictions
	<p>source database are consistent with those in the destination database.</p> <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. - If the value of session_replication_role of the destination database is replica when the full+incremental synchronization task is complete, change the value to the original one. <ul style="list-style-type: none"> ● Forcibly stop a task. <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is <code>drs_unique_ID</code>. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The lc_monetary values of the source and destination databases must be the same. ● To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● The destination database cannot contain objects with the same type and name as the objects to be synchronized, including databases, schemas, and tables. System databases, system schemas, and system tables are excluded. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. ● Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● For a full+incremental or incremental synchronization task, if an internal error occurs during the pre-check and the task stops before it is started, check and delete the streaming replication slot by referring to Forcibly Stopping

Type	Restrictions
	<p>Synchronization of PostgreSQL to avoid residual streaming replication slots in the source database.</p> <ul style="list-style-type: none"> ● After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. ● If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database. <p>NOTE DDL statements are captured using event triggers in the source database, recorded in a specific table, and then synchronized to the destination database. You need to create event triggers and functions in the source database before starting a task. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p> <ul style="list-style-type: none"> ● The restrictions on the objects that can be synchronized are as follows: <ul style="list-style-type: none"> - Object name: The database name cannot contain +"%\<>, the schema name and table name cannot contain ".\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). - Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. - Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. - Function: Do not synchronize C language functions or functions with the leakproof or support attribute. - Plug-in: The metadata of plug-ins is not synchronized. - Data type: Basic data types are not synchronized. - DRS does not synchronize BLOB data. - Type conversion: The binary coercion type cannot be converted. - Event trigger: Event triggers can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later. - Text search parser: Text search parsers can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later. - Text search template: Text search templates can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later.

Type	Restrictions
	<ul style="list-style-type: none"> - User: Existing users in the destination database, superuser, replication, and bypassrls attributes of users, and member relationships of superuser users are not synchronized. If the object owner or grantor is superuser, its owner or grantor is not synchronized. If the destination database is Huawei Cloud RDS for PostgreSQL DB instance, the password of the user to be synchronized cannot contain the username. During table-level synchronization, the default access permissions of source database users are not synchronized. After being synchronized to the destination database, the superuser user becomes a common user and is granted the root permissions by default (supported only when the destination database version is RDS for PostgreSQL 11.11 or later).

Procedure

This section uses to-the-cloud synchronization from PostgreSQL to PostgreSQL as an example to describe how to configure a real-time synchronization task in the VPC network scenario.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-77 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0256

Table 3-90 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-78 Synchronization instance information

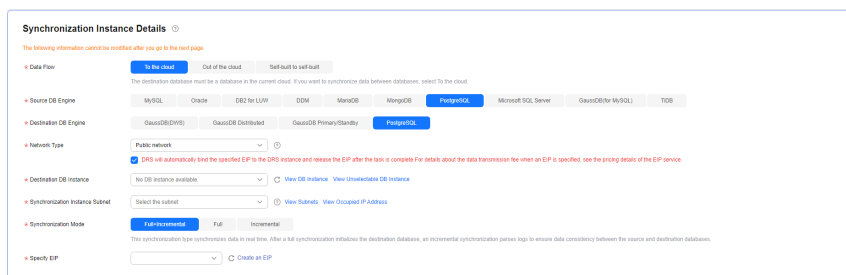


Table 3-91 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select PostgreSQL .
Network Type	Available options: VPC, Public network and VPN or Direct Connect . VPC is used as an example. <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
Destination DB Instance	The RDS for PostgreSQL DB instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-79 Task type




Table 3-92 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-80 Enterprise Project and Tags




Table 3-93 Enterprise Project and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

 **NOTE**

The source database can be an ECS database or an RDS instance. Configure parameters based on the database type.

- Scenario 1: Databases on an ECS - source database configuration

Figure 3-81 Self-build on ECS - source database information

Source Database
System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

[Test Connection](#) This button is available only after the replication instance is created successfully.

Table 3-94 Self-build on ECS - source database information

Parameter	Description
Database Type	Select Self-built on ECS .
VPC	A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Scenario 2: RDS DB instance - source database configuration

Figure 3-82 RDS DB instance - source database information

Source Database

Database Type: self-built database **RDS DB Instance**

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-95 RDS DB instance - source database information

Parameter	Description
Database Type	Select an RDS DB instance.
DB Instance Name	Select the RDS PostgreSQL instance to be synchronized as the source DB instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Figure 3-83 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

✔ Test successful

Table 3-96 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL instance you selected when creating the migration task and cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization objects and accounts and click **Next**.

Figure 3-84 Synchronization Object

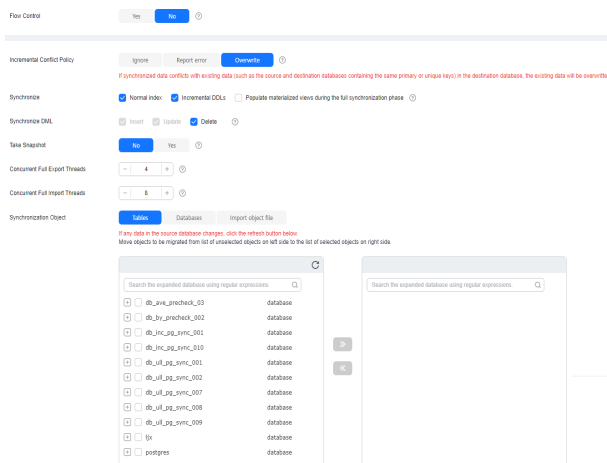
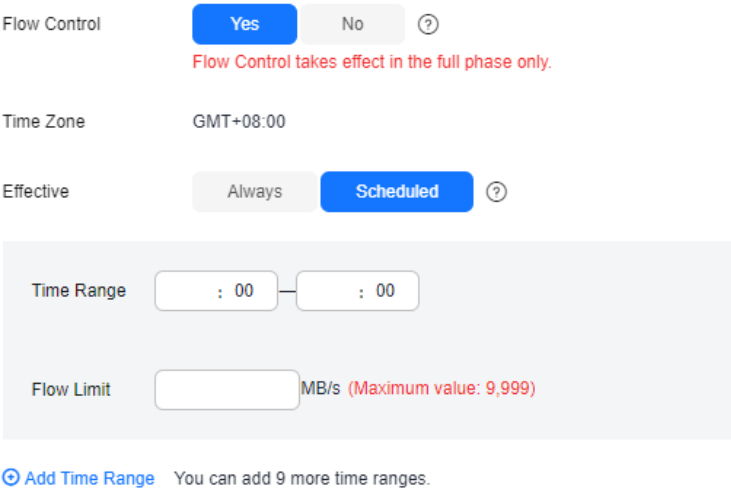


Table 3-97 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-85 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database.● Report error The synchronization task will be stopped and fail.● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize	<p>Available options: Normal index, Incremental DDLs, and Populate materialized views during the full synchronization phase</p> <p>Populate materialized views during the full synchronization phase: This option takes effect only for materialized views that was populated in the source database. This operation affects the full synchronization performance. You perform this operation after the full synchronization is complete.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Take Snapshot	<p>If you perform a full+incremental migration, you can take a snapshot for your databases.</p> <p>Exporting data in snapshot mode in the full export phase can effectively improve the data synchronization efficiency in the full+incremental export scenario. However, the snapshot mechanism of PostgreSQL prevents historical data in the database from being reclaimed during the export, which may cause space expansion. You are advised to use this method when the full or incremental data volume is large and the source database disk space is sufficient.</p>
Concurrent Full Export Threads	Number of threads for exporting objects and data during full synchronization.
Concurrent Full Import Threads	Number of threads for importing objects and data during full synchronization.


Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • Database-level synchronization: In full synchronization, the selected databases and the inventory data of the database objects are synchronized. In incremental synchronization, the DML and some DDL statements of all tables except unlogged tables and temporary tables are synchronized. • Table-level synchronization: In full synchronization, the inventory data of the selected tables, sequences, views, or materialized views is synchronized. In incremental synchronization, the DML and some DDL statements of the selected tables are synchronized. • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, after the synchronization, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.
Synchronize Account	<p>During the synchronization, you can synchronize accounts based on your service requirements. For details, see Table 3-98.</p>

Figure 3-86 Synchronize Account

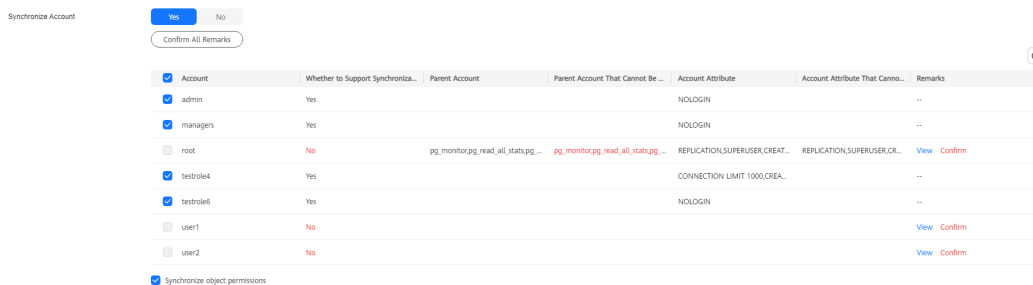


Table 3-98 Accounts and permissions to be synchronized

Parameter	Description
Account	Account name of the source database.
Whether to Support Synchronization	Whether the account can be synchronized. There are accounts that can be synchronized and accounts that cannot be synchronized. For an account that cannot be synchronized, the specific reason is displayed in View in the Remarks column.
Parent Account	Parent account.
Parent Account That Cannot Be Synchronized	The parent account that cannot be synchronized.
Account Attribute	Attributes of the source database account.
Account Attribute That Cannot Be Synchronized	The account attributes that cannot be synchronized due to insufficient permissions of the destination database user.
Remarks	Description of the parent account and account attributes that cannot be synchronized. You can go to the next step only after confirming all remarks.
Synchronize object permissions	Whether to synchronize permissions corresponding to the account.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-87 Task startup settings

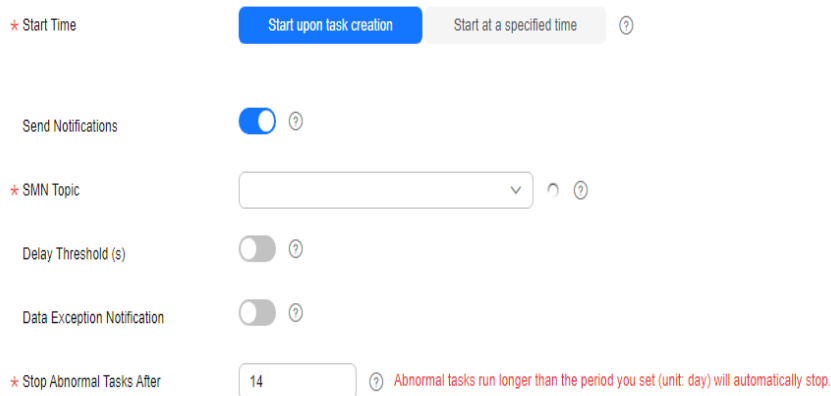



Table 3-99 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.9 From PostgreSQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-100 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14)ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14)Other cloud database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14)RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14)	GaussDB(DWS) clusters (versions 8.1.3 and 8.2.0)

 **NOTE**

Only whitelisted users can use this function.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-101](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-101 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	Database CONNECT permission, schema USAGE permission, table SELECT permission, and sequence SELECT permission	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permissions for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> 1. Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes.
Destination database user	<ul style="list-style-type: none"> • Database-level: The CREATEDB permission is required. • Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the objects are required. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-102](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-102 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. • Supported field types: <ul style="list-style-type: none"> - GaussDB(DWS) (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, real, double, bit, bit varying, boolean, bytea, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz, interval, cidr, path, box, lseg, macaddr, point, polygon, inet, tsquery, tsvector, uuid, json and jsonb, are supported. - XML, line, and domain synchronization is not supported. • Scope of full synchronization <ul style="list-style-type: none"> - Schemas, tables, indexes, constraints, sequences, and user-defined types During table-level synchronization, only tables and sequences can be synchronized. - Not supported: system schemas (any schema starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random), and system catalogs. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> • Object name: The database name cannot contain +"%\<>. The schema name and table name cannot contain ".'\<>. The column name cannot contain ". The column name cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. • Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. • Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> • Scope of incremental synchronization <ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables

Type	Precautions
	<ul style="list-style-type: none"> - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes <code>pg_logical_emit_message</code> to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-103](#).

Table 3-103 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The partition table trigger of the source database cannot be set to disable. • To perform incremental synchronization: The <code>pg_hba.conf</code> file of the source database contains the following configuration: <code>host replication all 0.0.0.0/0 md5</code> • The <code>test_decoding</code> plug-in must be installed on the source database in advance.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <ul style="list-style-type: none"> ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Some DML statements, including INSERT, UPDATE, and DELETE, are supported. ● DDL statements are not supported. ● DML statements of unlogged tables and temporary tables are not supported. ● DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. <p>Stopping a task</p> <ul style="list-style-type: none"> ● Stop a task normally. When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted.

Type	Restrictions
	<ul style="list-style-type: none"> ● Forcibly stop a task. <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The block_size value of the destination database must be greater than that of the source database. ● The destination database cannot contain objects with the same type and name as the objects to be synchronized, including databases, schemas, and tables. System databases, system schemas, and system tables are excluded. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● Only the primary table can be synchronized. The primary table will be converted into a common table and synchronized to the destination database. Data in the partition table will be written to the primary table. ● After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original constraint name</i> ● After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> ● Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails.

Type	Restrictions
	<ul style="list-style-type: none"> ● After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. ● In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. ● The restrictions on the objects that can be synchronized are as follows: <ul style="list-style-type: none"> - Object name: The database name cannot contain +'%\<>. The schema name and table name cannot contain ".!<>. The column name cannot contain ". The column name cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. - Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. - Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized.

Procedure

This section uses PostgreSQL to GaussDB(DWS) as an example to describe how to use DRS to configure a real-time synchronization task in the VPC network scenario.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-88 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name:

Description:

0/256

Table 3-104 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-89 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

• Data Flow: To the cloud (Selected) Out of the cloud Self-back to self-built

• Source DB Engine: MySQL Oracle DB2 for LUW IBM MariaDB Microsoft SQL Server PostgreSQL (Selected) GaussDB for MySQL TDSE

• Destination DB Engine: GaussDB(DWS) (Selected) GaussDB Distributed GaussDB Primary/Standby PostgreSQL

• Network Type: Public network Private network

• Creation DB Instance: No DB instance available (Selected) View DB Instance View Unavailable DB Instance

• Synchronization Instance Subnet:

• Synchronization Mode: Full synchronization (Selected)

• Specify EP:

Table 3-105 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB(DWS) .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. VPC is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The destination is a GaussDB(DWS) instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing. Ensure that the table structure of the selected object has been created in the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 3-90 Task type




Table 3-106 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-91 Enterprise Project and Tags




Table 3-107 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

 **NOTE**

The source database can be an ECS database or an RDS instance. Configure parameters based on different scenarios.

- Scenario 1: Databases on an ECS - source database configuration

Figure 3-92 Self-built on ECS - source database information

Source Database
System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

[Test Connection](#) This button is available only after the replication instance is created successfully.

Table 3-108 Self-built on ECS - source database information

Parameter	Description
Source Database Type	Select Self-built on ECS .
VPC	A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 - 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	SSL encrypts the connections between the source and destination databases.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Scenario 2: RDS DB instance - source database configuration

Figure 3-93 RDS DB instance - source database information

Source Database

DRS migrates only some key parameters to the destination database. For the other parameters that cannot be migrated, you need to use parameter templates to configure them on the destination database.

Database Type: self-built database RDS DB instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-109 RDS DB instance - source database information

Parameter	Description
Source Database Type	Select RDS DB instance .
DB Instance Name	Select the RDS PostgreSQL instance to be synchronized as the source DB instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Figure 3-94 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

✔ Test successful

Table 3-110 Destination database information

Parameter	Description
DB Instance Name	The GaussDB(DWS) instance selected when you created the migration task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization objects and accounts and click **Next**.

Figure 3-95 Synchronization Mode

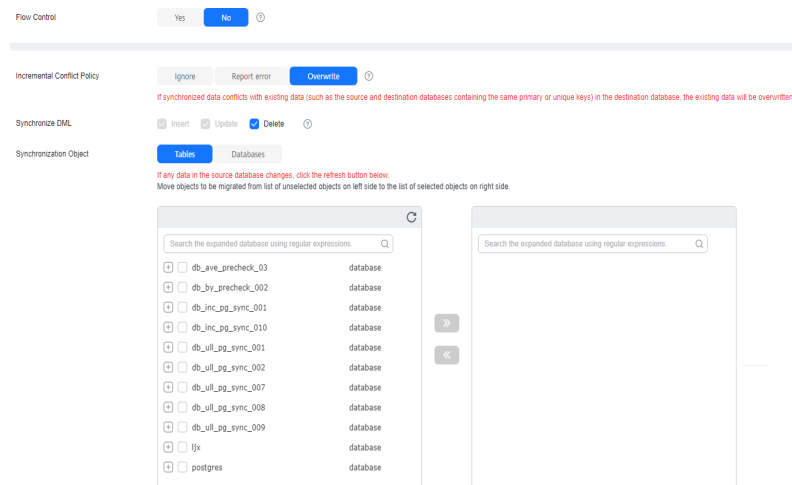
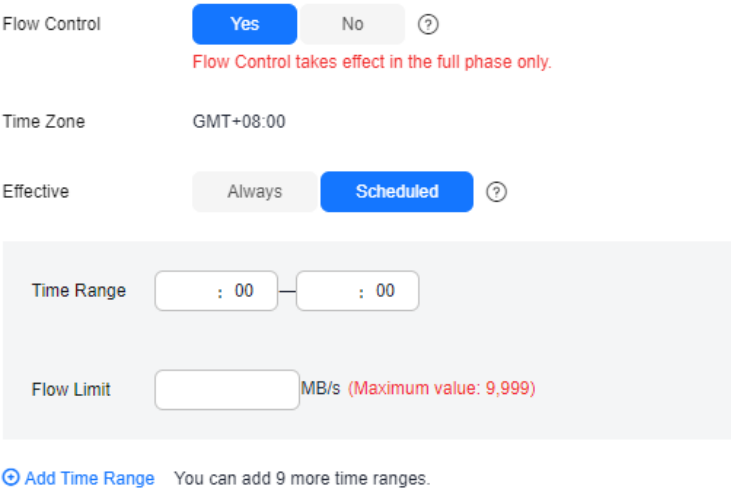



Table 3-111 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-96 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-97 Task startup settings

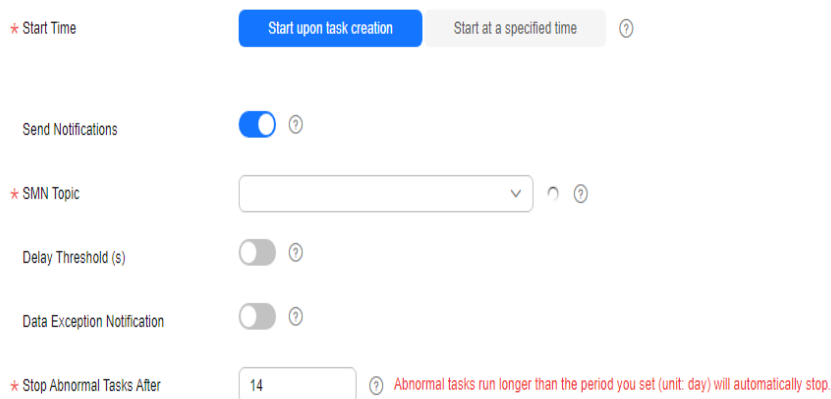



Table 3-112 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.10 From PostgreSQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-113 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) Other cloud database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14) 	GaussDB Primary/Standby (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-114](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-114 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full Synchronization	Full+Incremental Synchronization
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-115](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-115 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> - GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, real, double, bit, bit varying, boolean, bytea, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz, interval, cidr, path, box, lseg, macaddr, point, polygon, inet, tsquery, tsvector, uuid, json and jsonb, are supported. - XML, line, domain, and self-built data type synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> - Supported: schemas, tables, primary key and unique constraints, table data, and sequences - Not supported: system schemas and system catalogs (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain + "%\<>", the schema name and table name cannot contain ".\<>", and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. ● Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes <code>pg_logical_emit_message</code> to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-116](#).

Table 3-116 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The partition table trigger of the source database cannot be set to disable. • To perform incremental synchronization: <ul style="list-style-type: none"> - The pg_hba.conf file of the source database contains the following configuration: <pre>host replication all 0.0.0.0/0 md5</pre> - The <code>test_decoding</code> plug-in must be installed on the source database in advance.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <ul style="list-style-type: none"> ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Some DML statements, including INSERT, UPDATE, and DELETE, are supported. ● DDL statements are not supported. ● DML statements of unlogged tables and temporary tables are not supported. ● DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE</p> <p>Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes <code>pg_logical_emit_message</code> to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p> <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> • Do not limit the synchronization speed during data comparison. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The max_prepared_transactions value of the destination database must be greater than that of the source database. ● The max_worker_processes value of the destination database must be greater than that of the source database. ● The max_locks_per_transaction value of the destination database must be greater than that of the source database. ● The max_connections value of the destination database must be greater than that of the source database. ● The lc_monetary value of the destination database must be the same as that of the source database. ● The character set of the destination database must be the same as that of the source database. ● The time zone of the destination database must be the same as that of the source database. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● Only the primary table can be synchronized. The primary table will be converted into a common table and synchronized to the destination database. Data in the partition table will be written to the primary table. ● When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original

Type	Restrictions
	<p>index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key.</p> <ul style="list-style-type: none"> ● Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. ● After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. ● In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. ● During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. ● Case conversion is not performed on names of objects after they are synchronized to the destination database. ● The restrictions on the objects that can be synchronized are as follows: <ul style="list-style-type: none"> – Object name: The database name cannot contain +"%\<>, the schema name and table name cannot contain ".!<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). – Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. – Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-98 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Region] .. v
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Project] .. v

* Task Name: DRS-5678 ⓘ

Description: [Description] ⓘ
0/256

Table 3-117 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-99 Synchronization instance details

Synchronization Instance Details ⓘ

⚠ The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** (Out of the cloud) | Self-built to self-built

• Source DB Engine: MySQL | Oracle | DR2 for LUW | DOM | MariaDB | MongoDB | **PostgreSQL** | Microsoft SQL Server | GaussDB for MySQL | TDS

• Destination DB Engine: GaussDB(DRS) | GaussDB Distributed | **GaussDB Priority/Standby** | PostgreSQL

• Network Type: Public network ⓘ

• Destination DB Instance: No DB instance available | View DB Instance | View OracleDB DB Instance

• Synchronization Instance Subnet: Select the subnet | View Subnets | View Occupied IP Address

• Synchronization Mode: **Full Synchronization** | Full

• Specify EIP: [EIP] | Check an EIP

Table 3-118 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . VPC is used as an example. <ul style="list-style-type: none">– VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.– Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.– VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB primary/standby instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full This synchronization mode is suitable for scenarios where service interruption is acceptable. Database objects and data, including tables, views, and stored procedures, from non-system databases can be synchronized to the destination all at once. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-100 AZ




Table 3-119 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-101 Enterprise Project and Tags




Table 3-120 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

 **NOTE**

The source database can be an ECS database or an RDS instance. Configure parameters based on different scenarios.

- Scenario 1: Databases on an ECS - source database configuration

Figure 3-102 Self-build on ECS - source database information

Source Database
System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

This button is available only after the replication instance is created successfully.

Table 3-121 Self-build on ECS - source database information

Parameter	Description
Database Type	Select Self-built on ECS .
VPC	A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Scenario 2: RDS DB instance - source database configuration

Figure 3-103 RDS DB instance - source database information

Source Database

DRS migrates only some key parameters to the destination database. For the other parameters that cannot be migrated, you need to use parameter templates to configure them on the destination database.

Database Type: self-built database RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

SSL Connection:

Table 3-122 RDS DB instance - source database information

Parameter	Description
Source Database Type	Select an RDS DB instance.
DB Instance Name	Select the RDS PostgreSQL instance to be synchronized as the source DB instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Figure 3-104 Destination database information

Destination Database

DB Instance Name

Database Username


Database Password 

Table 3-123 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-105 Synchronization mode

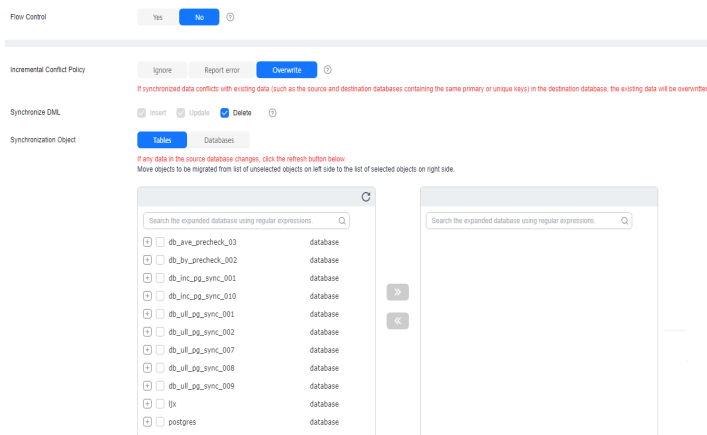
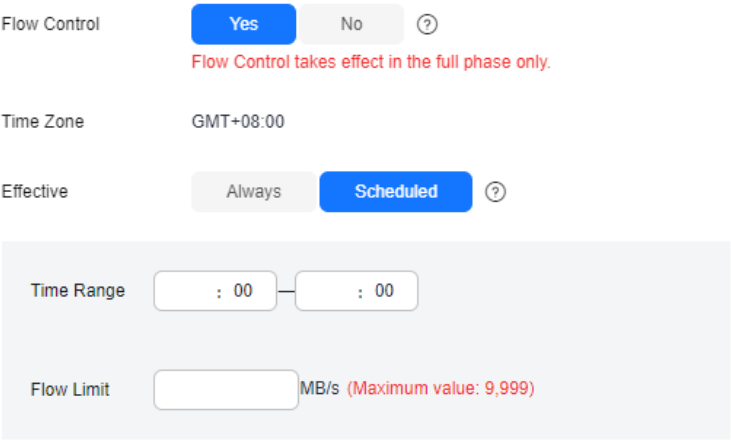



Table 3-124 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-106 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-107 Task startup settings





Table 3-125 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.11 From PostgreSQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-126 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) Other cloud database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14) 	GaussDB Distributed (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-127](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-127 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full Synchronization	Full+Incremental Synchronization
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-128](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-128 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> – GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, real, double, bit, bit varying, boolean, bytea, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestamptz, interval, cidr, path, box, lseg, macaddr, point, polygon, inet, tsquery, tsvector, uuid, json and jsonb, are supported. – XML, line, domain, and self-built data type synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> – Supported: schemas, tables, primary key and unique constraints, table data, and sequences – Not supported: system schemas and system catalogs (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). – A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail.

Type	Precautions
	<p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> • Object name: The database name cannot contain +'%'\<>, the schema name and table name cannot contain ".'\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). • Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. • Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> • Scope of incremental synchronization <ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-129](#).

Table 3-129 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">● The partition table trigger of the source database cannot be set to disable.● To perform incremental synchronization:<ul style="list-style-type: none">- The pg_hba.conf file of the source database contains the following configuration: host replication all 0.0.0.0/0 md5- The test_decoding plug-in must be installed on the source database in advance.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <ul style="list-style-type: none"> • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Some DML statements, including INSERT, UPDATE, and DELETE, are supported. • DDL statements are not supported. • DML statements of unlogged tables and temporary tables are not supported. • DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE</p> <p>Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes <code>pg_logical_emit_message</code> to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p> <p>Synchronization comparison</p> <ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> • Do not limit the synchronization speed during data comparison. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The max_prepared_transactions value of the destination database must be greater than that of the source database. ● The max_worker_processes value of the destination database must be greater than that of the source database. ● The max_locks_per_transaction value of the destination database must be greater than that of the source database. ● The max_connections value of the destination database must be greater than that of the source database. ● The lc_monetary value of the destination database must be the same as that of the source database. ● The character set of the destination database must be the same as that of the source database. ● The time zone of the destination database must be the same as that of the source database. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● Only the primary table can be synchronized. The primary table will be converted into a common table and synchronized to the destination database. Data in the partition table will be written to the primary table. ● When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original

Type	Restrictions
	<p>index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: c_+hash value+original constraint name (which may be truncated)+_key.</p> <ul style="list-style-type: none"> • After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. • Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. • In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. • Case conversion is not performed on names of objects after they are synchronized to the destination database. • If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. • The restrictions on the objects that can be synchronized are as follows: <ul style="list-style-type: none"> – Object name: The database name cannot contain +"%\<>, the schema name and table name cannot contain ".\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). – Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in

Type	Restrictions
	<p>the primary table can be synchronized. All data in the partition table will be written to the primary table.</p> <ul style="list-style-type: none"> - Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-108 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-130 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-109 Synchronization instance details

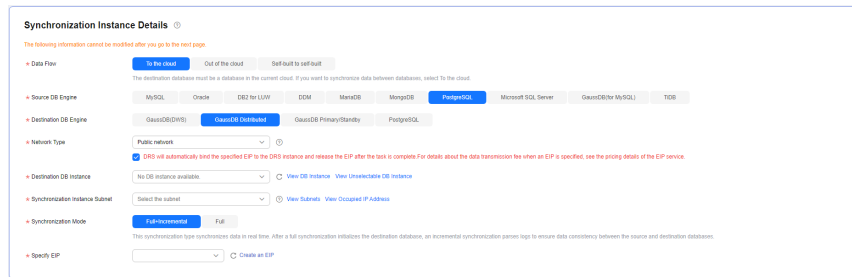


Table 3-131 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB Distributed .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . VPC is used as an example. <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB distributed instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full This synchronization mode is suitable for scenarios where service interruption is acceptable. Database objects and data, including tables, views, and stored procedures, from non-system databases can be synchronized to the destination all at once. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-110 AZ




Table 3-132 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-111 Enterprise Project and Tags




Table 3-133 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

 **NOTE**

The source database can be an ECS database or an RDS instance. Configure parameters based on different scenarios.

- Scenario 1: Databases on an ECS - source database configuration

Figure 3-112 Self-build on ECS - source database information

Source Database
System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

This button is available only after the replication instance is created successfully.

Table 3-134 Self-build on ECS - source database information

Parameter	Description
Database Type	Select Self-built on ECS .
VPC	A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Scenario 2: RDS DB instance - source database configuration

Figure 3-113 RDS DB instance - source database information

Source Database

Database Type: self-built database RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-135 RDS DB instance - source database information

Parameter	Description
Database Type	Select an RDS DB instance.
DB Instance Name	Select the RDS PostgreSQL instance to be synchronized as the source DB instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Figure 3-114 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-136 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-115 Synchronization mode

Flow Control: Yes No

Incremental Conflict Policy: Ignore Report error Overwrite

Synchronize DML: Insert Update Delete

Synchronization Object: Tables Databases

If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

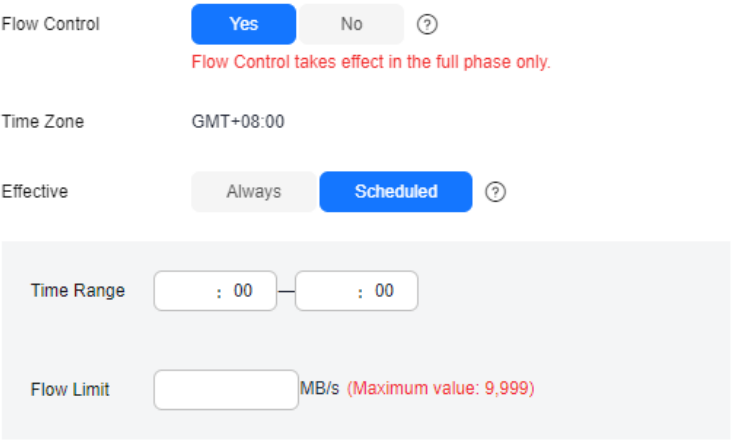
Search the expanded database using regular expressions:

- db_save_precheck_03 database
- db_mv_precheck_002 database
- db_mv_pg_sync_001 database
- db_mv_pg_sync_010 database
- db_mv_pg_sync_001 database
- db_mv_pg_sync_002 database
- db_mv_pg_sync_007 database
- db_mv_pg_sync_008 database
- db_mv_pg_sync_009 database
- ltx database
- postgres database


Search the expanded database using regular expressions:

>> <<

Table 3-137 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-116 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-117 Task startup settings

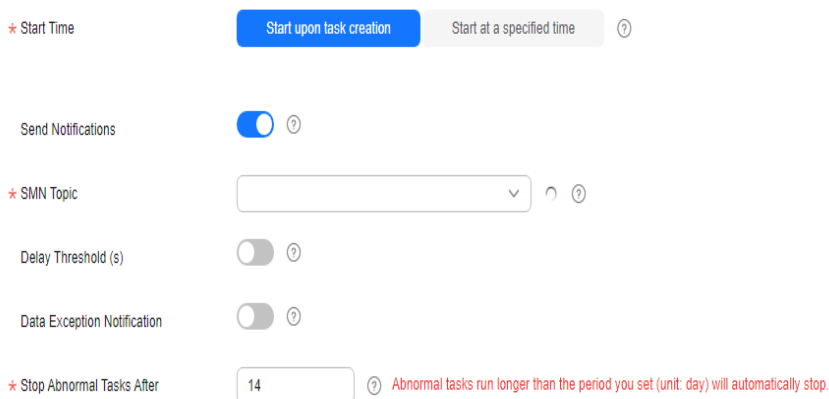



Table 3-138 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.12 From Oracle to MySQL

Supported Source and Destination Databases

Table 3-139 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) 	RDS for MySQL (5.5, 5.6, 5.7, and 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization

tasks require different permissions. For details, see [Table 3-140](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-140 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table, EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER permissions for a CDB. Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; • Oracle 12c or later in non-tenant mode: <p>You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <p>Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;</p> • To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; <ul style="list-style-type: none"> • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>SESSION, SELECT ANY DICTIONARY, SELECT for a single table, EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;</p>
Destination database user	The user must have the SELECT, INSERT, CREATE, DROP, UPDATE, ALTER, DELETE and INDEX permissions.	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-141 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-141 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● Table-level synchronization or object file import is supported. ● Databases, table structures, primary keys, unique keys, common indexes, and table data can be synchronized. ● Other database objects, such as stored procedures, triggers, functions, sequences, packages, synonyms, or users, cannot be synchronized. ● In the full synchronization phase, bfile, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. ● Partitions in the table structure cannot be synchronized. Partitioned tables are changed to non-partitioned tables after being synchronized to the destination database. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● An empty source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-142](#).

Table 3-142 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. • The maximum row length of Oracle cannot exceed 8 KB, excluding BLOB and TEXT columns which will be automatically converted to the text and blob types of MySQL. The reason is that the MySQL InnoDB restricts the row length to 8 KB. • The primary key or unique key column cannot contain values of the string data type when you map the MySQL data types to the character data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur. • The values of binary_float and binary_double cannot be set to Nan, Inf, or -Inf because MySQL does not support these values. DRS converts the three values to 0 and saves them by default. • MySQL does not support the synchronization of the check constraints of Oracle. • AUTO_PK_ROW_ID cannot be used as a column name in Oracle because it is a reserved column name in MySQL 5.7 and cannot be created. • Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle. • The precision of the float(p) field in the Oracle database cannot exceed the precision range p: [1, 32]. The float range of MySQL is smaller than that of Oracle. • The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. • The total index length of columns in the source database cannot exceed the length limit in the destination database.

Type	Restrictions
	<p>For detailed length requirements, see Index Length Description.</p> <ul style="list-style-type: none">• The Default User statement is not supported in MySQL.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. Do not perform operations (including but not limited to DDL and DML operations) on the destination database. Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not change the char field in the source database or destination database table to varchar, or it is padded with extra spaces due to differences between Oracle and MySQL. In this case, data inconsistency may occur. During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met.

Type	Restrictions
	<p>If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database.</p> <ul style="list-style-type: none"> - During table-level synchronization, basic DDLs (ALTER TABLE ADD COLUMN, ALTER TABLE DROP COLUMN, ALTER TABLE RENAME COLUMN, ALTER TABLE MODIFY COLUMN, and TRUNCATE TABLE) are supported. Default values cannot be modified. - Database-level synchronization supports CREATE TABLE. (Table definitions cannot contain functions.) - The object in a DDL cannot be the keyword of the destination database, such as index or where. For details about keywords of the destination MySQL database, see MySQL official documentation. - If the destination database version is earlier than 8.0, alter table rename column is not supported. - Incremental DDL operations do not support special characters such as full-width and Chinese characters. • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. • If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. • After tables are synchronized to the destination database, their names are converted into lowercase letters. For example, the name of table ABC is changed to abc after being synchronized to the destination database. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail.

Type	Restrictions
	<ul style="list-style-type: none">• When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● During synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. ● The time zone settings of the source and destination database must be the same. ● When you select to synchronize the table structure, the destination instance cannot contain the database to be synchronized. ● The storage of the destination database should be about 1.5 times greater than the storage of the source database. ● If the destination database version is earlier than 5.7.7, the index column length cannot exceed 767 bytes. If the destination database version is later than 5.7.7, the length cannot exceed 3072 bytes. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code.

Type	Restrictions
	<ul style="list-style-type: none"> • After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4. • If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. • If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. • In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. • For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. • If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. • There are some syntax differences between Oracle and MySQL, so the syntax including but not limited to functions, expressions, and referenced system tables, may not be completely converted during the structure synchronization. Therefore, during the synchronization, the structure may exist in the Oracle database but does not exist in the MySQL database, or the syntax exists in the MySQL database but is not converted. As a result, the structure fails to be synchronized. If this happens, you can manually create a table structure in the destination database. • You can enable strong consistency during synchronization to ensure that the commit sequence and atomicity of transactions synchronized to the destination database are the same as those of the source database. However, the performance in this mode is much lower than that in the default mode. • When strong consistency is enabled, out-of-memory (OOM) may occur if the size of committed transactions is greater than 256 MB. • In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB.

Type	Restrictions
	<ul style="list-style-type: none"> ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. ● When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by primary or foreign keys ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. ● Only Oracle 11g and later versions support XMLTYPE data synchronization. ● In the Oracle database, if the default value is followed by a comment starting with --, a space must be added after --. Otherwise, a syntax error will be reported when the value is migrated to the MySQL database because the MySQL database requires that -- be followed by a space.

Procedure

This section uses real-time synchronization from Oracle to RDS for MySQL as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-118 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page. The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name:

Description:

0/256

Table 3-143 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-119 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud Out of the cloud Self-built to self-built

Source DB Engine: MySQL Oracle DB2 for LUW DB2 MingDB WingDB PostgreSQL Microsoft SQL Server GaussDB for MySQL TDSE

Destination DB Engine: MySQL DB2 GaussDB(DWS) GaussDB Distributed GaussDB Primary/Standby PostgreSQL GaussDB for MySQL

Network Type: Public network

DRS will automatically bind the specified EP to the DRS instance and unbind the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing detail of the EP service.

Destination DB Instance: No DB instance available View DB Instance View Unavailable DB Instance

During the full synchronization of a DRS task, a lot of storage are generated. These storage may be non-replicable cloud disks, which may cause the storage space to be used up. You are advised to enable storage autoresizing for the RDS DB instance. During the DRS task, set an appropriate auto-resizing period for RDS groups. You can also clear storage according the specified retention period with just a few clicks.

Synchronization Instance Subnet: Select the subnet New Subnets View Occupied IP Address

Synchronization Mode: Full/Incremental Full Incremental

Create Binlog Cleanup:

Specify EP: Create an EP

Table 3-144 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select MySQL .
Network Type	Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-120 Task type




Table 3-145 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-121 Enterprise Project and Tags




Table 3-146 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-122 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name
For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

PDB Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 3-147 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-123 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-148 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MySQL instance selected when you created the migration task. The instance cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted. You can change the password if necessary.
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-124 Synchronization mode

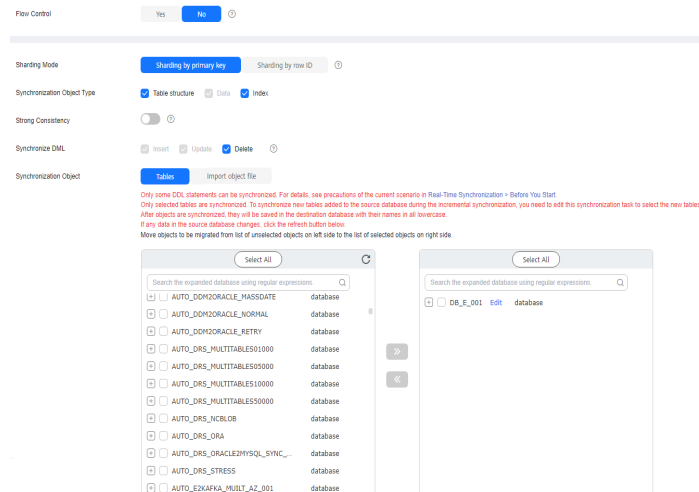
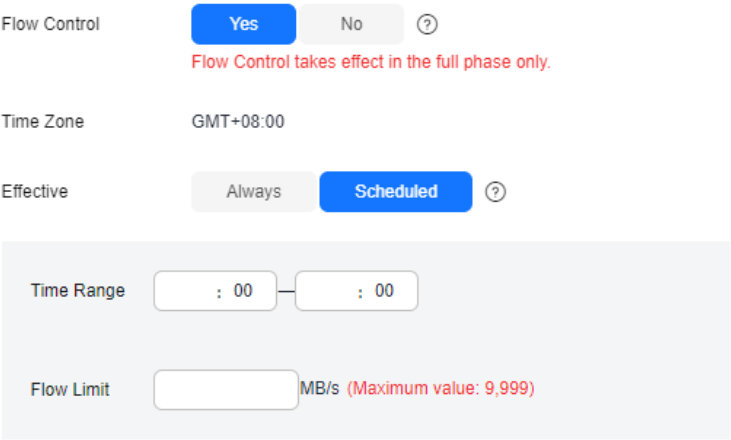



Table 3-149 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-125 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Strong Consistency	<p>This option is disabled by default. You can enable it to ensure that the commit sequence and atomicity of transactions synchronized to the destination database are the same as those of the source database. However, the performance in this mode is much lower than that in the default mode.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required. To quickly select the desired database objects, you can use the search function.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Advanced Settings** page, set the parameters for the incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 3-126 Advanced settings

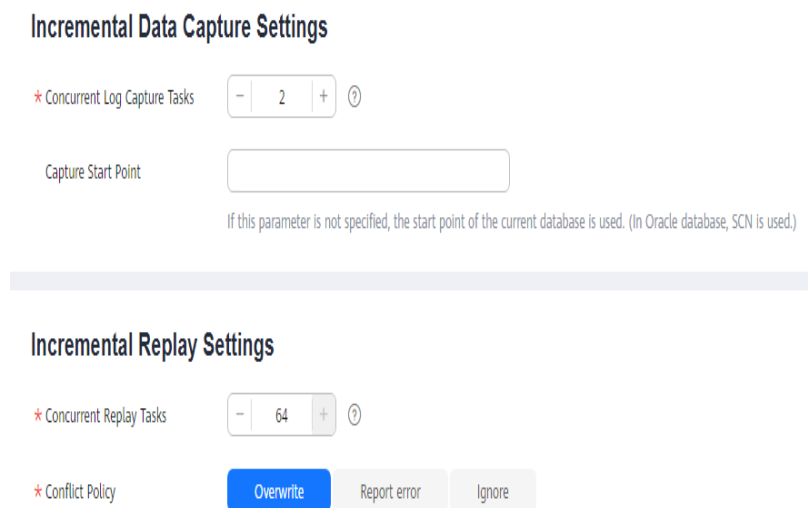


Table 3-150 Incremental capture settings

Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads that read logs from the source database. The value ranges from 1 to 16. Each thread reads logs in the sequence of log files.	2
Capture Start Point	Specifies the SCN for starting the capture. SCNs are designed to meet service requirements. It consists of a start SCN for capturing and a valid SCN. For details, see the SCN concepts of Oracle. For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?	If this parameter is left blank, the current SCN of the database is used as the start point by default.

Table 3-151 Incremental replay settings

Parameter	Description	Default Value
Concurrent Replay Tasks	The number of concurrent threads for writing data to the destination database. The value ranges from 1 to 64.	64
Conflict Policy	<ul style="list-style-type: none"> Overwrite The data captured by DRS will overwrite the data in the destination database. Report error An error message is displayed, indicating that the synchronization task is abnormal. Ignore The system skips the error record and continues the data replay. 	Overwrite

Step 6 On the **Process Data** page, filter the data to be synchronized and click **Next**. For details, see [Processing Data](#).

Step 7 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 8 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-127 Task startup settings

Table 3-152 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 9 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.13 From Oracle to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-153 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	GaussDB(for MySQL) Primary/Standby

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-154](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-154 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and CREATE 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization
	<p>SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; <ul style="list-style-type: none"> • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p> <ul style="list-style-type: none"> • Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.
Destination database user	The user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD, CREATE VIEW, CREATE ROUTINE and TRIGGER.	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-155 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-155 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Databases, table structures, primary keys, unique keys, common indexes, and table data can be synchronized. • Other database objects, such as stored procedures, triggers, functions, sequences, packages, synonyms, or users, cannot be synchronized. • In the full synchronization phase, bfile, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. • In the incremental synchronization phase, bfile, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • An empty source database cannot be synchronized. • Tables with virtual columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-156](#).

Table 3-156 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> ● The maximum row length of Oracle cannot exceed 8 KB, excluding BLOB and TEXT columns which will be automatically converted to the text and blob types of MySQL. The reason is that the MySQL InnoDB restricts the row length to 8 KB. ● The primary key or unique key column cannot contain values of the string data type when you map the MySQL data types to the character data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur. ● The values of binary_float and binary_double cannot be set to Nan, Inf, or -Inf because MySQL does not support these values. DRS converts the three values to 0 and saves them by default. ● MySQL does not support the synchronization of the check constraints of Oracle. ● AUTO_PK_ROW_ID cannot be used as a column name in Oracle because it is a reserved column name in MySQL 5.7 and cannot be created. ● Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. ● The total index length of columns in the source database cannot exceed the length limit in the destination database. For detailed length requirements, see Index Length Description. ● The Default User statement is not supported in MySQL. ● An empty source database cannot be synchronized.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. • Do not change the char field in the source database or destination database table to varchar, or it is padded with extra spaces due to differences between Oracle and MySQL. In this case, data inconsistency may occur. <p>Full synchronization</p> <ul style="list-style-type: none"> • In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename

Type	Restrictions
	<p>column, alter table modify column, and truncate table. The modification of default values is not supported.</p> <ul style="list-style-type: none"> - Database-level synchronization supports create table. (Table definitions cannot contain functions.) - The object in a DDL cannot be the keyword of the destination database, such as index or where. - Incremental DDL operations do not support special characters such as full-width and Chinese characters. <ul style="list-style-type: none"> ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● After tables are synchronized to the destination database, their names are converted into lowercase letters. For example, the name of table ABC is changed to abc after being synchronized to the destination database. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> • When you select to synchronize the table structure, the destination instance cannot contain the database to be synchronized. • During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. • Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. • The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • Associated objects must be synchronized at the same time to avoid synchronization failure due to missing associated objects. Common dependencies: tables referenced by primary or foreign keys • The time zone settings of the source and destination database must be the same. • The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. • During synchronization, writing data to the destination databases is not allowed. Otherwise, data inconsistency may occur. • If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database.

Type	Restrictions
	<ul style="list-style-type: none"> ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. ● After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4. ● If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database.

Type	Restrictions
	<ul style="list-style-type: none"> ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● The names of mapped databases and tables are case-insensitive, which means no matter if the object name is uppercase or lowercase, it stays lowercase after the object is synchronized to the destination database. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. ● When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. ● The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. ● Only Oracle 11g and later versions support XMLTYPE data synchronization. ● In the Oracle database, if the default value is followed by a comment starting with --, a space must be added after --. Otherwise, a syntax error will be reported when the value is migrated to the MySQL database because the MySQL database requires that -- be followed by a space.

Procedure

This section uses real-time synchronization from Oracle to GaussDB(for MySQL) as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-128 Synchronization task information

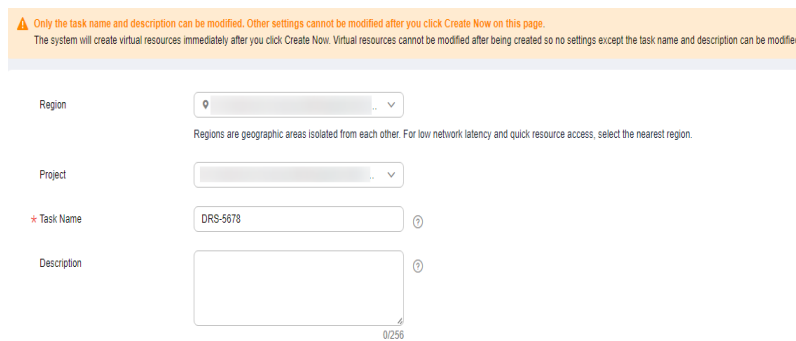


Table 3-157 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-129 Synchronization instance details

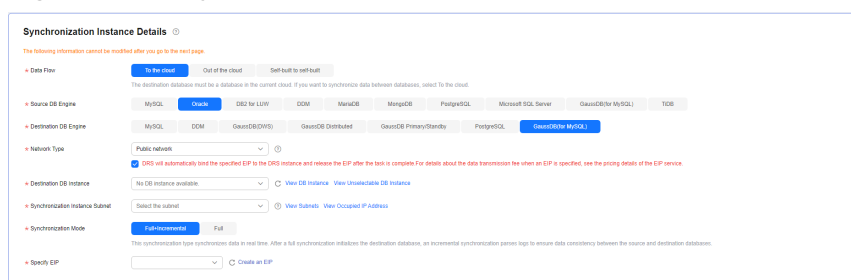


Table 3-158 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The GaussDB(for MySQL) instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-130 Task type




Table 3-159 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-131 Enterprise Project and Tags




Table 3-160 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-132 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

PDB Name ⓘ

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 3-161 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-133 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-162 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted. You can change the password if necessary.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-134 Synchronization mode

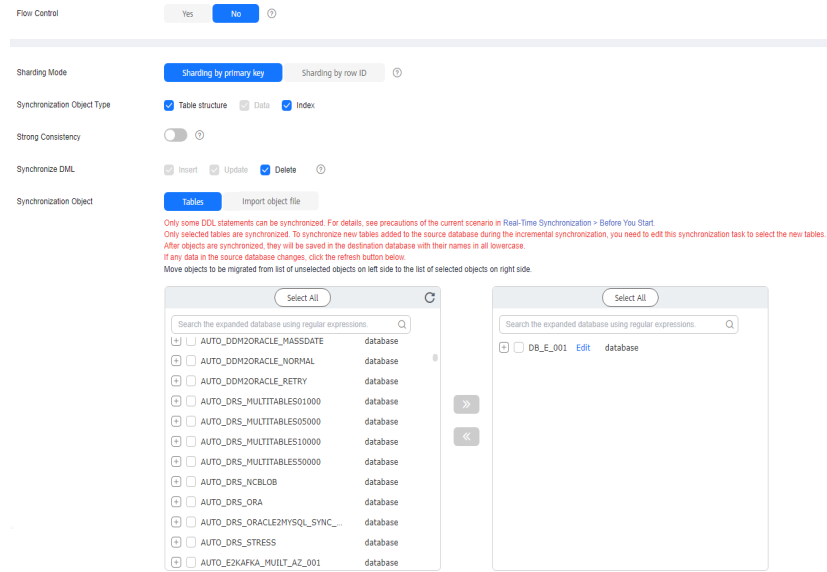
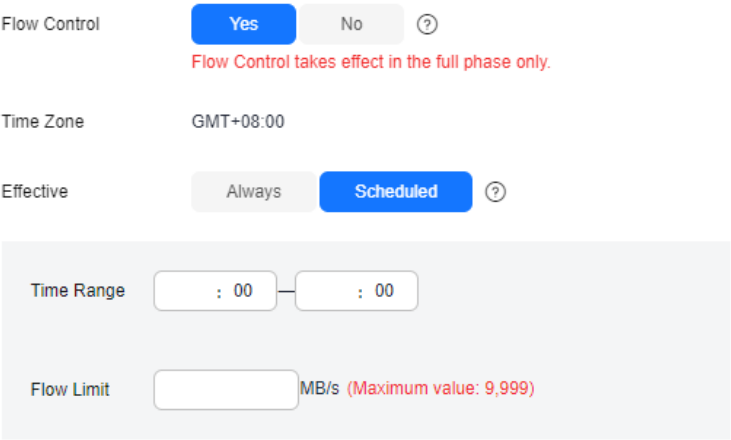



Table 3-163 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-135 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.


Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p>
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>


Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required. To quickly select the desired database objects, you can use the search function.</p> <ul style="list-style-type: none"> For details about how to import an object file, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Advanced Settings** page, set the parameters for the full synchronization or full+incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 3-136 Full synchronization

Full Synchronization Settings

* Concurrent Task Configuration 

* Concurrent Export Tasks 


* Concurrent Import Tasks 

Table 3-164 Full synchronization settings

Parameter	Description	Default Value
Concurrent Task Configuration	Indicates whether to configure the number of concurrent import and export tasks.	Enable
Concurrent Export Tasks	The number of export threads. The value ranges from 1 to 64. A larger value indicates higher load on the source database.	8

Parameter	Description	Default Value
Concurrent Import Tasks	The number of import threads. In stream mode, the number of concurrent import tasks must be equal to the number of concurrent export tasks. In non-stream mode, there is no such restriction. The value ranges from 1 to 64. A larger value indicates higher load on the destination database.	8

Figure 3-137 Incremental synchronization

Incremental Data Capture Settings

* Concurrent Log Capture Tasks ?

Table 3-165 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads that read logs from the source database. The value ranges from 1 to 16. Each thread reads logs in the sequence of log files.	2

Step 6 On the **Process Data** page, filter the data to be synchronized and click **Next**. For details, see [Processing Data](#).

Step 7 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 8 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-138 Task startup settings





Table 3-166 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 9 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.14 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-167 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) 	GaussDB Primary/Standby (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization

tasks require different permissions. For details, see [Table 3-168](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-168 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; <ul style="list-style-type: none"> • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user		<ul style="list-style-type: none"> • Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> • Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> • Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> • If <code>gs_loader</code> is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-169](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-169 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. • Full synchronization does not support the following column types: bfile, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, sdo_geometry, urowid, interval, and user-defined types. • The maximum precision supported by timestamp and interval day to second is 6. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. • For incremental synchronization, the LOB type supports only the BasicFiles attribute and does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. • For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • Tables with virtual columns in the source database cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-170](#).

Table 3-170 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"><li data-bbox="587 342 1385 472">• The total index length of columns in the source database cannot exceed the length limit in the destination database. For detailed length requirements, see Index Length Description.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> • Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. • Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. • During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize

Type	Restrictions
	<p>the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During table structure synchronization in a full synchronization, only default value constraints of the character string and number types are supported. Default value constraints of the function and sequence types are not supported. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. - During incremental synchronization, tables cannot be renamed, tables cannot be redefined online, and partitions cannot be swapped. ● During incremental synchronization, modify the start point of a capture task to resynchronize data. <ul style="list-style-type: none"> - After the capture point is modified, the previous object-level comparison result is not displayed. - If the start point of a capture task is changed separately, the change is synchronized to the start point of the replay task. That is, the start point of the replay task is the same as that of the capture task. This does not affect the change of the start point of the replay task. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not

Type	Restrictions
	<p>generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency.</p> <ul style="list-style-type: none"> ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: decode, nvl, nvl2, instr, substr, to_char, to_date, to_timestamp, length, lengthb, sysdate, trunc, nullif, next_day, regexp_instr, add_months, systimestamp, to_number, empty_clob and empty_blob. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● By default, the Oracle-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If

Type	Restrictions
	<p data-bbox="667 297 1422 427">newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> <li data-bbox="592 443 1366 539">● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. <li data-bbox="592 555 1414 752">● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data types supported by the primary key, see Mapping Data Types. <li data-bbox="592 768 1426 898">● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. <li data-bbox="592 913 1426 1077">● If the Oracle character set is WE8MSWIN1252 or WE8ISO8859P1, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. <li data-bbox="592 1093 1426 1189">● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. <li data-bbox="592 1205 1426 1335">● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. <li data-bbox="592 1350 1426 1447">● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. <li data-bbox="592 1462 1426 1592">● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. <li data-bbox="592 1608 1426 1872">● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. <li data-bbox="592 1888 1426 1984">● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. ● If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the destination database sequence maximum value plus the

Type	Restrictions
	<p>security margin, and the auto-decrement sequence value is the destination database sequence minimum value minus the security margin. The default security margin is 10,000.</p> <ul style="list-style-type: none"> • If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization. • In the Oracle database, if the default value is followed by a comment starting with --, a space must be added after --. Otherwise, a syntax error will be reported when the value is migrated to the MySQL database because the MySQL database requires that -- be followed by a space.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-139 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-171 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-140 Synchronization instance details

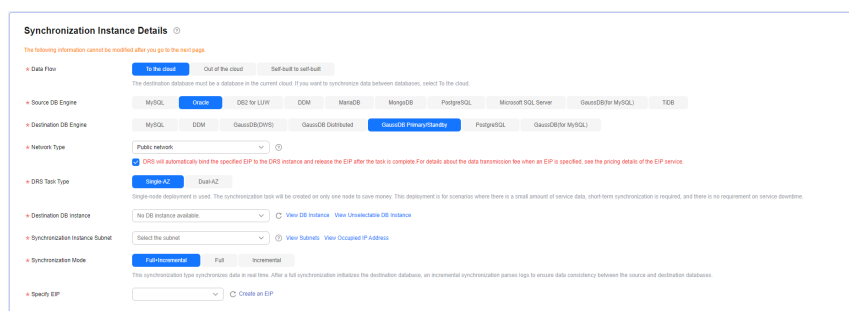


Table 3-172 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB Primary/Standby .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	The GaussDB primary/standby instance you created.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles .
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. If DRS Task Type is set to Dual-AZ , you need to specify the primary and standby IP addresses. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 3-141 Task type





Table 3-173 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-142 AZ</p> 

- Enterprise Project and Tags

Figure 3-143 Enterprise Project and Tags




Table 3-174 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 3-144 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

PDB Name ⓘ

Database Username

Database Password ⓘ

SSL Connection ⓘ

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

This button is available only after the replication instance is created successfully.

Table 3-175 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Figure 3-145 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-176 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.

Parameter	Description
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 3-146 Synchronization mode

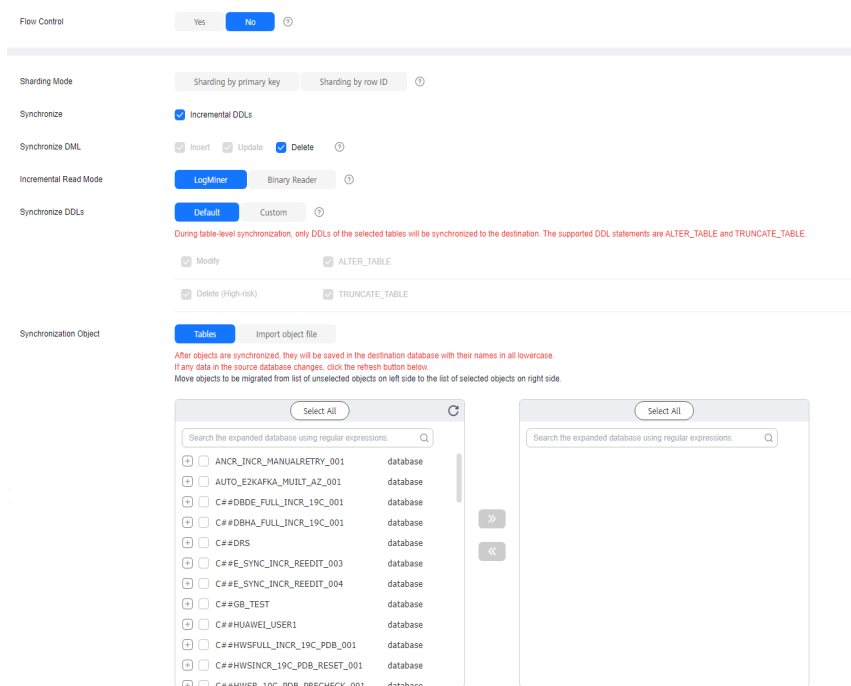
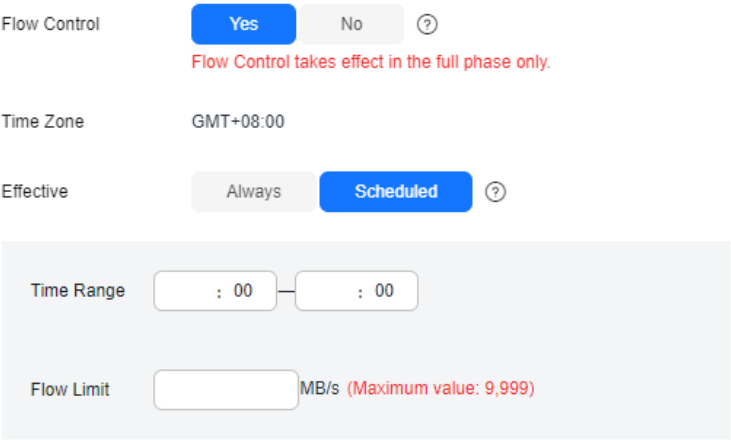



Table 3-177 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-147 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

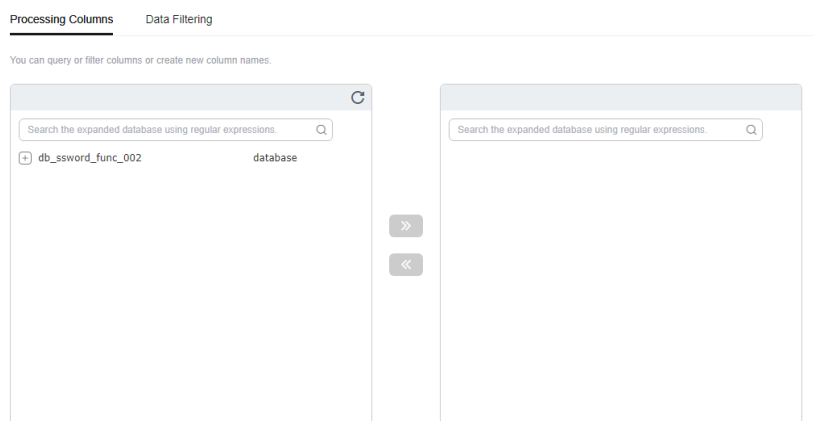
Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p>
Synchronize	Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Incremental Read Mode	<p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • Binary Reader: uses DRS-developed method to directly read and parse original redo logs. The performance is high, and the Oracle resource consumption is low. <p>LogMiner is recommended.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 3-148 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-149 Task startup settings

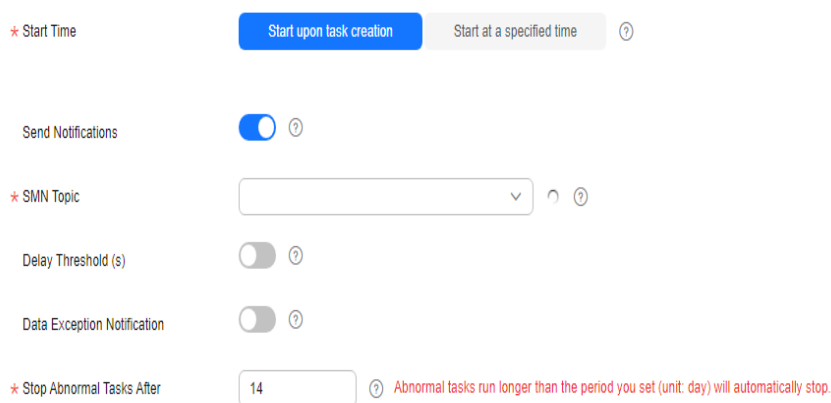



Table 3-178 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.15 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 3-179 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	GaussDB Distributed (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-180](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-180 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user		<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-181](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-181 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● Full synchronization does not support the following column types: bfile, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, sdo_geometry, urowid, interval year to month, interval day to second, and user-defined types. ● The maximum precision supported by timestamp and interval day to second is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-182](#).

Table 3-182 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"><li data-bbox="587 344 1433 479">• The total index length of columns in the source database cannot exceed the length limit in the destination database. For detailed length requirements, see Index Length Description.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> • Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. • Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. • During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize

Type	Restrictions
	<p>the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During table structure synchronization in a full synchronization, only default value constraints of the character string and number types are supported. Default value constraints of the function and sequence types are not supported. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. - During incremental synchronization, tables cannot be renamed, tables cannot be redefined online, and partitions cannot be swapped. ● During incremental synchronization, modify the start point of a capture task to resynchronize data. <ul style="list-style-type: none"> - After the capture point is modified, the previous object-level comparison result is not displayed. - If the start point of a capture task is changed separately, the change is synchronized to the start point of the replay task. That is, the start point of the replay task is the same as that of the capture task. This does not affect the change of the start point of the replay task. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not

Type	Restrictions
	<p>generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency.</p> <ul style="list-style-type: none"> ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● During incremental synchronization, if the source or destination database is abnormal, the task fails. After the database is recovered, the task is globally started. The status of the original capture or replay component is ignored, and the replay component is started from the capture interruption point. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: decode, nvl, nvl2, instr, substr, to_char, to_date, to_timestamp, length, lengthb, sysdate, trunc, nullif, next_day, regexp_instr, add_months, systimestamp, to_number, empty_clob and empty_blob. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● By default, the Oracle-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If

Type	Restrictions
	<p data-bbox="667 297 1422 427">newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> <li data-bbox="592 443 1366 539">● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. <li data-bbox="592 555 1414 752">● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data types supported by the primary key, see Mapping Data Types. <li data-bbox="592 768 1426 898">● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. <li data-bbox="592 913 1426 1077">● If the Oracle character set is WE8MSWIN1252 or WE8ISO8859P1, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. <li data-bbox="592 1093 1426 1189">● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. <li data-bbox="592 1205 1426 1335">● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. <li data-bbox="592 1350 1426 1447">● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. <li data-bbox="592 1462 1426 1592">● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. <li data-bbox="592 1608 1426 1872">● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. <li data-bbox="592 1888 1426 2013">● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the destination database is a distributed GaussDB database, the update operation performed on the source distribution column will fail to be synchronized to the GaussDB database during incremental synchronization, causing data inconsistency. Therefore, avoid updating a distribution column. ● You are advised to disable the global secondary index (GSI) function for the destination database. Otherwise, incremental synchronization may fail. ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● When editing the task to add a new table, ensure that transactions of the new table have been committed.

Type	Restrictions
	<p>Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours.</p> <ul style="list-style-type: none"> • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. • If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-150 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 3-183 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-151 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** (Out of the cloud, Self-built to self-built)

• Source DB Engine: MySQL (Oracle, DB2 for LUW, DDM, MaxDB, MongoDB, PostgreSQL, Microsoft SQL Server, GaussDB for MySQL, TDS)

• Destination DB Engine: MySQL, DDM, GaussDB(DWS), **GaussDB Distributed**, GaussDB Primary/Standby, PostgreSQL, GaussDB for MySQL

• Network Type: Public network (Private network)

• DRD Task Type: Single AZ (Dual-AZ)

• Destination DB Instance: No DB instance available (View DB Instance, View Unavailable DB Instance)

• Synchronization Instance Subnet: Select the subnet (View Subnets, View Occupied IP Address)

• Synchronization Mode: **Full-incremental** (Full, Incremental)

• Specify EP: (Create an EP)

Table 3-184 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .

Parameter	Description
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	The GaussDB distributed instance you created.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles .
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. If DRS Task Type is set to Dual-AZ , you need to specify the primary and standby IP addresses. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 3-152 Task type





Table 3-185 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-153 AZ</p> 

- Enterprise Project and Tags

Figure 3-154 Enterprise Project and Tags




Table 3-186 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 3-155 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

PDB Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 3-187 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Figure 3-156 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-188 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.

Parameter	Description
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 3-157 Synchronization mode

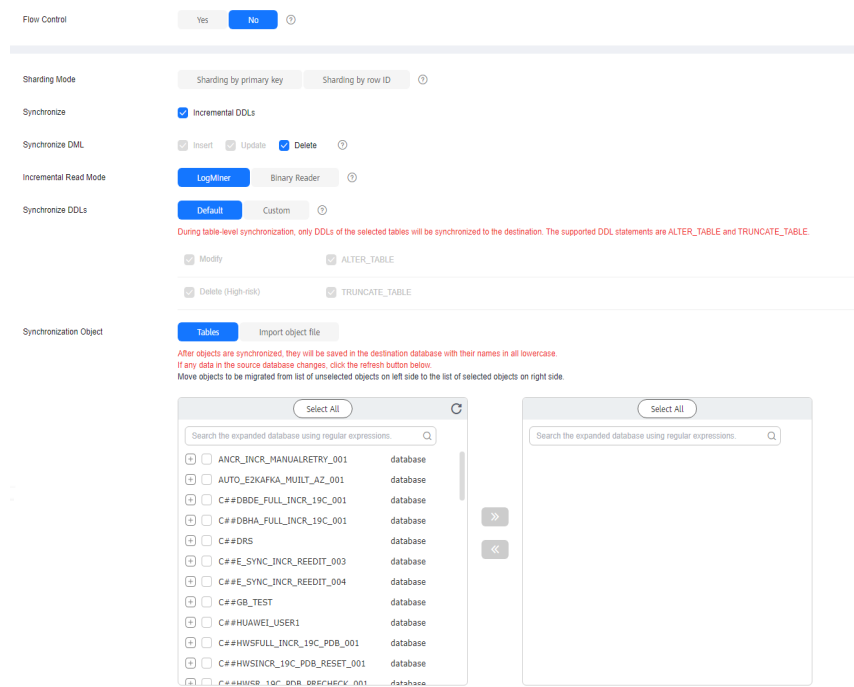
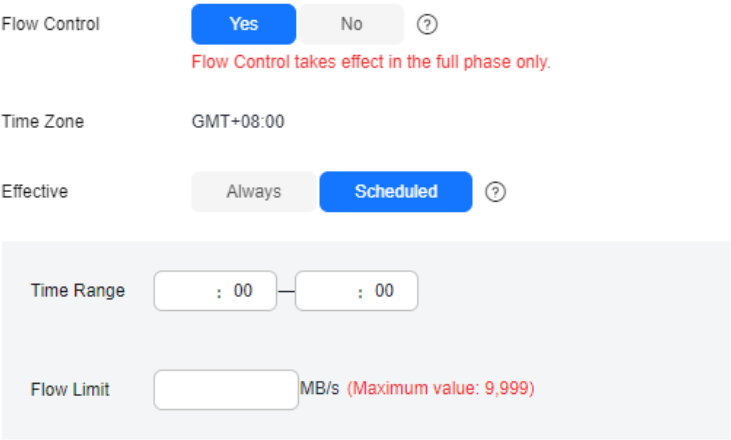



Table 3-189 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-158 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

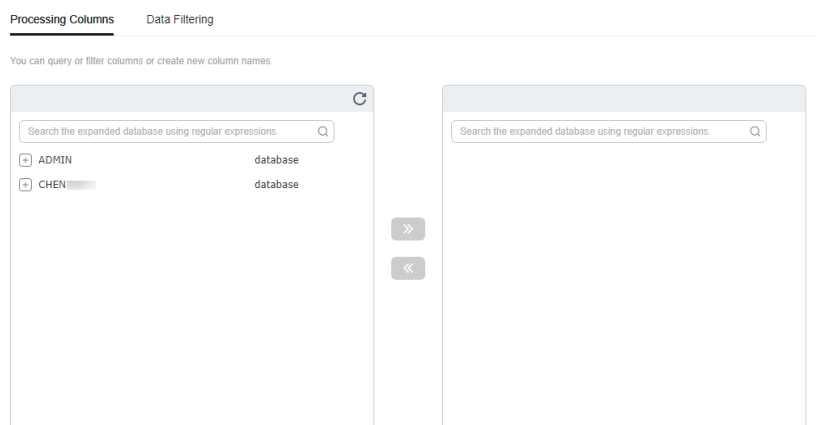
Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.
Synchronize	Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Incremental Read Mode	<p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • Binary Reader: uses DRS-developed method to directly read and parse original redo logs. The performance is high, and the Oracle resource consumption is low. <p>LogMiner is recommended.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 3-159 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-160 Task startup settings

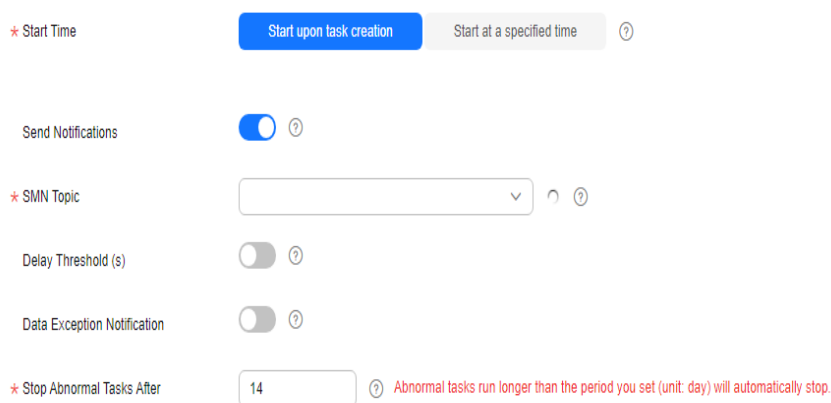



Table 3-190 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.16 From Oracle to DDM

Supported Source and Destination Databases

Table 3-191 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	DDM instances

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-192](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-192 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and CREATE 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization
	<p>SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; <ul style="list-style-type: none"> • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>
Destination database user	The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD and CREATE VIEW	

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-193 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-193 Supported synchronization objects

Type	Precautions
Synchronization objects	The source database data can be synchronized, but the source table structure and other objects cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-194](#).

Table 3-194 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The maximum row length of Oracle cannot exceed 8 KB, excluding BLOB and TEXT columns which will be automatically converted to the text and blob types of MySQL. The reason is that the MySQL InnoDB restricts the row length to 8 KB.• The primary key or unique key column cannot contain values of the string data type when you map the MySQL data types to the character data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur.• The values of binary_float and binary_double cannot be set to Nan, Inf, or -Inf because MySQL does not support these values. DRS converts the three values to 0 and saves them by default.• AUTO_PK_ROW_ID cannot be used as a column name in Oracle because it is a reserved column name in MySQL 5.7 and cannot be created.• Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle.• Tables whose names contain special characters except underscores (_) cannot be synchronized.• An empty source database cannot be synchronized.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • DDL operations cannot be performed on the source database. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. • During synchronization, writing data to the destination databases is not allowed. Otherwise, data inconsistency may occur. <p>Full synchronization</p> <ul style="list-style-type: none"> • In the full synchronization phase, bfile, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. • In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • DDL operations cannot be synchronized during incremental synchronization. • In the incremental synchronization phase, bfile, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported.

Type	Restrictions
	<ul style="list-style-type: none"> ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● You need to create table structures and indexes in the destination database based on table structure of the source schema. Objects that are not created in the destination database are not to be synchronized. ● During synchronization, create a table structure in the destination database. The table structure of the destination database must contain all columns in the source database, and the primary keys must be the same. ● The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> – The default value may be left empty. – The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● Before synchronization, you must create databases, tables, columns, indexes, and constraints with the same names (in lowercase) as the corresponding source objects to be synchronized in the destination database. ● During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables

Type	Restrictions
	<p>may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.</p> <ul style="list-style-type: none"> ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After objects such as tables are synchronized to the destination database, their names are converted into lowercase letters. For example, the name of object ABC is changed to abc after being synchronized to the destination database. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. ● After the Oracle table structure is synchronized to DDM, the character set of the table is utf8mb4. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. ● In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During an incremental synchronization of table-level objects, renaming tables is not recommended. ● The source database username is consistent with the destination database name. ● When the source database user and table structure are synchronized to the destination database, the corresponding information is converted into lowercase letters. For example, tables Ab and AB are changed to table ab after being synchronized to the destination database. ● Index organized tables cannot be synchronized. ● Before a full synchronization task starts, if a transaction is not committed for a long time, data may be lost. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours.

Type	Restrictions
	<ul style="list-style-type: none"> • If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-161 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-195 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-162 Synchronization instance details

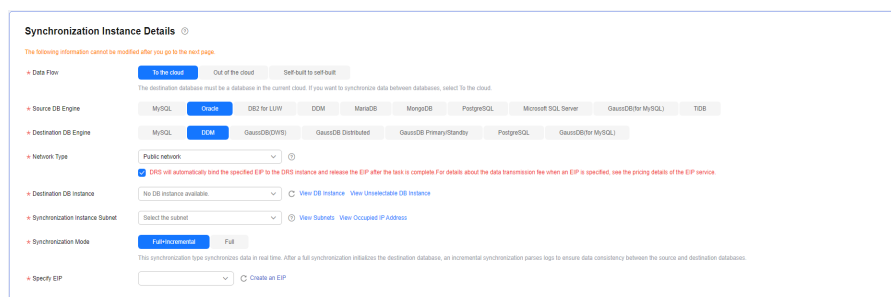


Table 3-196 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select DDM .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available DDM instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> – Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-163 Task type




Table 3-197 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-164 Enterprise Project and Tags




Table 3-198 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-165 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name
For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

PDB Name ⓘ

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 3-199 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-166 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-200 Destination database settings

Parameter	Description
DB Instance Name	The DDM instance you selected when you create a synchronization task. The instance name cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization object. Click **Next**.

Figure 3-167 Synchronization Mode

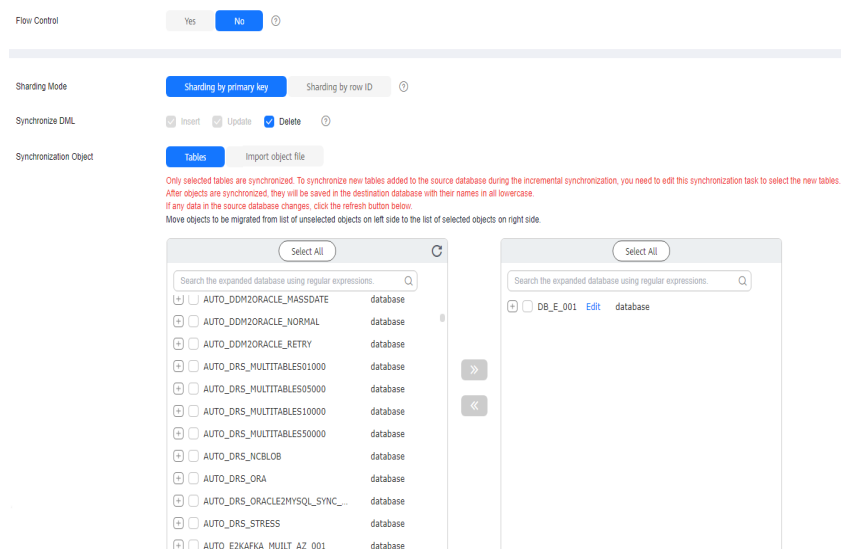
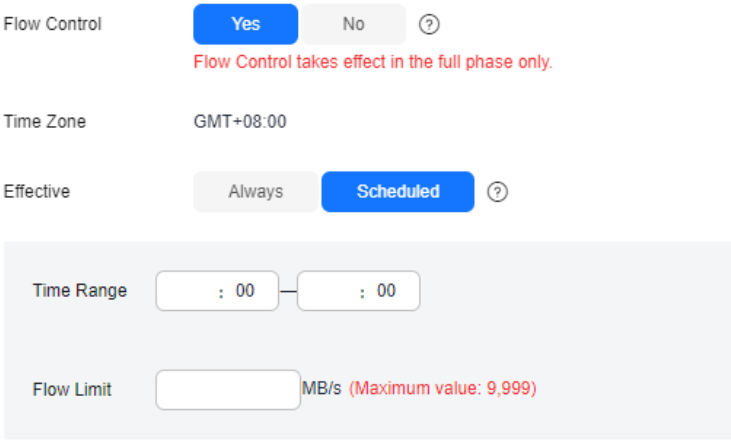



Table 3-201 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-168 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-169 Task startup settings

Table 3-202 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.17 From Oracle to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-203 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	GaussDB(DWS) clusters (versions 8.1.3 and 8.2.0)

 **NOTE**

Only whitelisted users can use this function.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-204](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-204 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; <ul style="list-style-type: none"> • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>
Destination database user	<ul style="list-style-type: none"> • If the destination side does not contain databases, schemas, or tables, the destination database user must have the permission to create databases, the permission to create schemas in a database, or the permission to create tables in a schema. • The INSERT, SELECT, UPDATE, and DELETE permissions are required for each table. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-205](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-205 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, common indexes (B-Tree indexes), and constraints (primary key, empty, and not-null) can be synchronized. • Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, unique indexes, foreign key indexes, and check constraints cannot be synchronized. • Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-206](#).

Table 3-206 Precautions

Type	Restrictions
Restrictions on the source database	An empty source database cannot be synchronized.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> • In the full synchronization phase, bfile, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. • In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • In the incremental synchronization phase, bfile, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. • In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. • Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated.

Type	Restrictions
	<ul style="list-style-type: none"> • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> – Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. ● Only the to_date and sys_guid functions can be used as default values. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following situations may occur: <ul style="list-style-type: none"> - The default value is left empty. - The table fails to be created. As a result, the task fails. - If DDLs for incremental synchronization contain system functions that are not found in the destination database, the synchronization task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Virtual columns in the source database table cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Materialized views in the source Oracle database will be converted into common tables after being synchronized to the destination database. ● The destination DB instance must have sufficient storage space. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables

Type	Restrictions
	<p>may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.</p> <ul style="list-style-type: none"> ● Associated objects must be synchronized at the same time to avoid synchronization failure due to missing associated objects. ● Object names, such as table names, are converted to lowercase letters after being synchronized to the destination database. For example, ABC is converted to abc. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail. ● If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● Only normal indexes are synchronized during index synchronization. Primary key constraints are synchronized with the table structure. ● If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index named a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● Multiple source tables can be mapped to one GaussDB(DWS) table. For details, see From Oracle to GaussDB(DWS). ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● The names of mapped databases and tables are case-insensitive, which means no matter if the object name is uppercase or lowercase, it stays lowercase after the object is synchronized to the destination database. ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database,

Type	Restrictions
	<p>the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All.</p> <ul style="list-style-type: none"> • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-170 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region .. v
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project .. v

* Task Name ⓘ

Description ⓘ

0/256

Table 3-207 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-171 Synchronization instance details

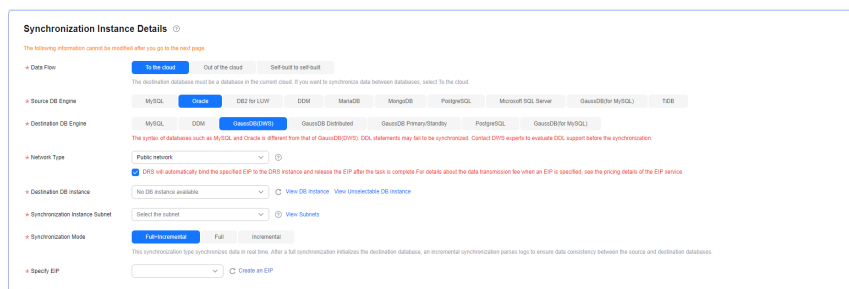


Table 3-208 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB(DWS) instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-172 Task type




Table 3-209 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-173 Enterprise Project and Tags




Table 3-210 Enterprise Project and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-174 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

PDB Name

Database Username

Database Password

SSL Connection

Table 3-211 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-175 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-212 Destination database settings

Parameter	Description
DB Instance Name	The default value is the GaussDB(DWS) instance selected for creating the migration task. It cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization object. Click **Next**.

Figure 3-176 Synchronization mode

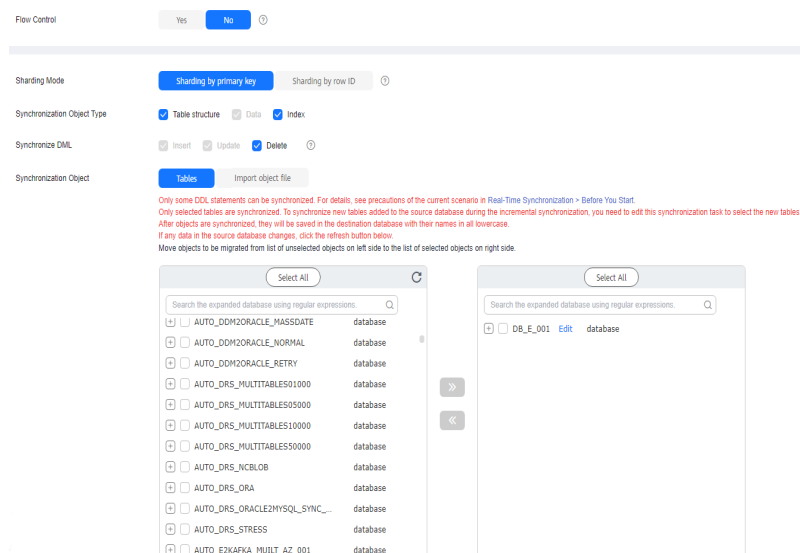
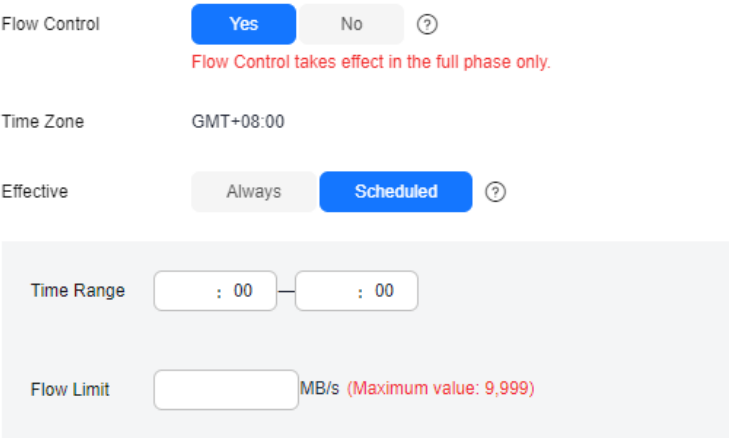



Table 3-213 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-177 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> ● Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. ● Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten.

Parameter	Description
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can synchronize tables or import object files based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Processing Data** page, filter data or add additional columns for the table object to be processed, and click **Next**.

- If you need to set data filtering, click **Data Filtering** and set related filtering rules.
- If you need to add additional columns, click the **Additional Columns** tab, click **Add** in the **Operation** column, and enter the column name and the operation type.

For details about related operations, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-178 Task startup settings

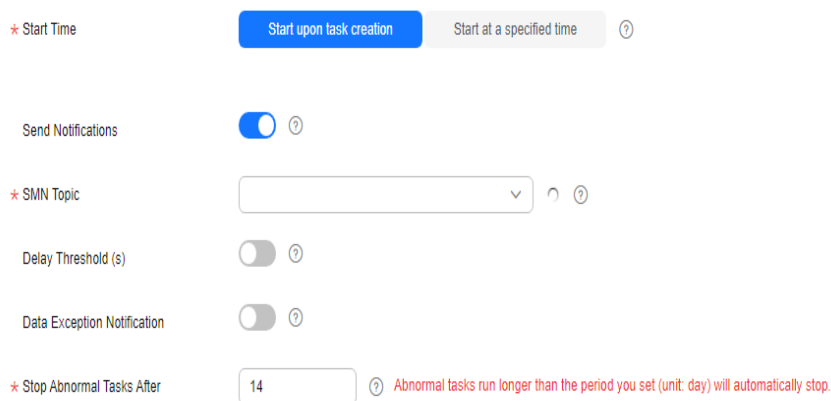



Table 3-214 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.18 From Oracle to PostgreSQL

Supported Source and Destination Databases

Table 3-215 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	RDS for PostgreSQL (versions 9.5, 9.6, 10, 11, 12, 13, and 14)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-216](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-216 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and CREATE 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization
	<p>SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>
Destination database user	The user must have the following permissions for a table: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-217 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-217 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Tables, indexes, and constraints (primary key, null, and non-null) can be synchronized.• Views, foreign keys, stored procedures, triggers, functions, events, and virtual columns cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-218](#).

Table 3-218 Precautions

Type	Restrictions
Restrictions on the source database	An empty source database cannot be synchronized.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. During a full synchronization, DRS writes large amount of data to the destination PostgreSQL database. As a result, the number of PostgreSQL WAL logs increases sharply, and the PostgreSQL disk space may be used up. You can disable the PostgreSQL log backup function before the full synchronization to reduce the number of WAL logs. After the synchronization is complete, enable the function. For details, see Setting an Automated Backup Policy. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p> <ul style="list-style-type: none"> During the synchronization, bfile, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. <p>Incremental synchronization</p>

Type	Restrictions
	<ul style="list-style-type: none"> • In the incremental synchronization phase, bfile, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. • In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. • Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database. Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. ● The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Restrictions
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • Associated objects must be synchronized at the same time to avoid synchronization failure due to missing associated objects. • Object names, such as table names, are converted to lowercase letters after being synchronized to the destination database. For example, ABC is converted to abc. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail. • If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. • If a task fails because the (self-built) table structure of a table without primary keys is not synchronized and the task cannot be resumed after the table structure is restored, reset the task. • During table structure synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. • If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. • If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. • If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component.

Type	Restrictions
	<ul style="list-style-type: none"> ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● Only normal indexes are synchronized during index synchronization. Primary key constraints are synchronized with the table structure. ● The names of mapped table-level objects are case-insensitive. For example, table ABC mapped to the destination database will be changed to table abc. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written

Type	Restrictions
	<p>to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All.</p> <ul style="list-style-type: none"> • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

This section uses real-time synchronization from Oracle to RDS for PostgreSQL as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-179 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-219 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-180 Synchronization instance details

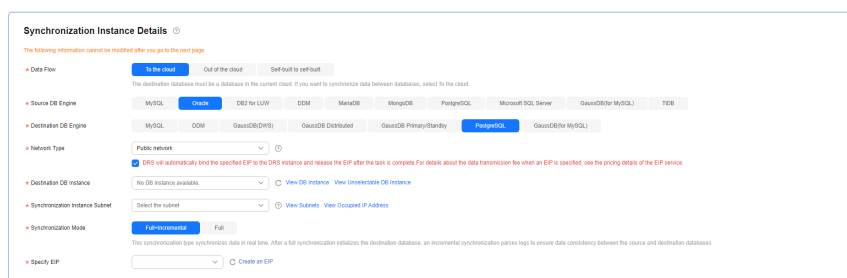


Table 3-220 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select PostgreSQL .
Network Type	Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The RDS for PostgreSQL instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> – Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-181 Task type




Table 3-221 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-182 Enterprise Project and Tags




Table 3-222 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-183 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name
For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

PDB Name ⓘ

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 3-223 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-184 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-224 Destination database settings

Parameter	Description
DB Instance Name	The RDS for PostgreSQL instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted. You can change the password if necessary.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-185 Synchronization mode

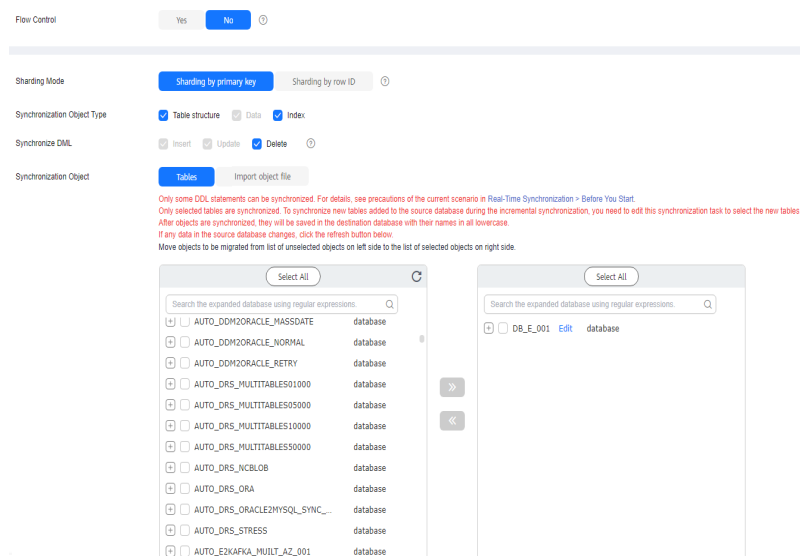
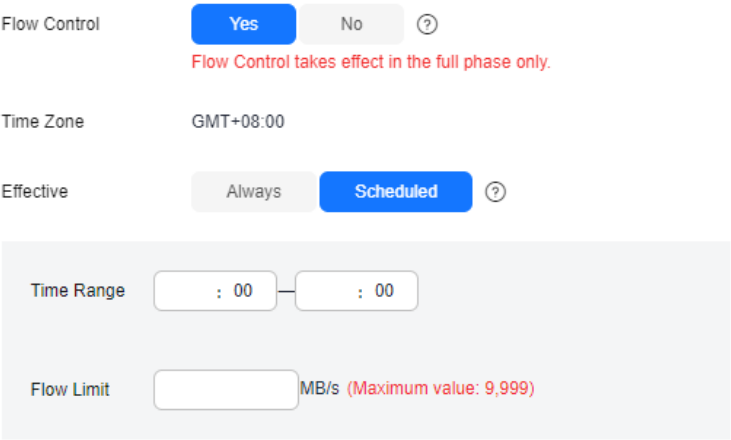



Table 3-225 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-186 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p>
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> Data is selected by default. If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can synchronize tables or import object files based on your service requirements. To quickly select the desired database objects, you can use the search function.</p> <ul style="list-style-type: none"> For details about how to import an object file, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-187 Task startup settings





Table 3-226 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.19 From DDM to MySQL

Supported Source and Destination Databases

Table 3-227 Supported databases

Source DB	Destination DB
DDM instances	RDS for MySQL (5.5, 5.6, 5.7, and 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-228](#). DRS automatically

checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-228 Database account permission

Type	Full, Full+Incremental, and Incremental
Source database user	<ul style="list-style-type: none"> The user of the source DDM database must have at least one permission, for example, SELECT. The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions.
Destination database user	<p>The user must have the SELECT, CREATE, ALTER, DROP, DELETE, INSERT and UPDATE permissions.</p> <p>The root account of the RDS for MySQL DB instance has the preceding permissions by default.</p>

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-229](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-229 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> Full synchronization supports the synchronization of data, table structures, and indexes. The source database cannot contain tables whose sharding keys are of the timestamp type. The sharding key of the source table must be added to the primary key and unique key of the destination table, which means that the primary key and unique key columns of the destination table must contain the sharded columns of the source table to avoid data conflict and inconsistency.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source

and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-230](#).

Table 3-230 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: .'<>\ • Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. <p>Full synchronization</p> <p>During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not perform data restoration on the source database. Some DDL operations are supported. <ul style="list-style-type: none"> DROP_DATABASE, DROP_TABLE, TRUNCATE_TABLE, CREATE_VIEW and DROP_VIEW are not supported. Online DDL is not supported. Tables can be created. For example: <pre>create table `ddl_test` (id int, c1 varchar(25), primary key(id)); create table `ddl_test_gho` like `ddl_test`;</pre> Tables can be renamed. Both the source and destination tables must be selected. For example: <pre>rename table `ddl_test` to `ddl_test_new`;</pre> Columns in a table can be added and modified, but cannot be deleted. For example: <pre>alter table `ddl_test` add column `c2` varchar(25); alter table `ddl_test` modify column `c1` varchar(50); alter table `ddl_test` alter c1 set default '***';</pre> Table indexes can be modified. For example: <pre>alter table `ddl_test` drop primary key; alter table `ddl_test` add primary key(id); alter table `ddl_test` add index `ddl_test_uk`(id); alter table `ddl_test` drop index `ddl_test_uk`;</pre> In table-level synchronization, you can add columns, modify columns, and add primary keys and normal indexes.

Type	Restrictions
	<ul style="list-style-type: none"> - During database-level synchronization, you can create tables, rename tables, add columns, modify columns, and add primary keys and normal indexes. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. • Perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. • The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● Resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down. ● If the source database contains a duplicate primary key or unique key, the data synchronized to the destination database will be less than that in the source database. Therefore, you must check and correct the data before starting the synchronization task. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If a table in the source database contains system columns <code>_ddm_lock</code> and <code>_slot</code>, data in these columns will be filtered out during full synchronization and incremental synchronization, resulting in data loss. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, the destination database cannot be set to read-only. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-188 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-231 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-189 Synchronization instance information

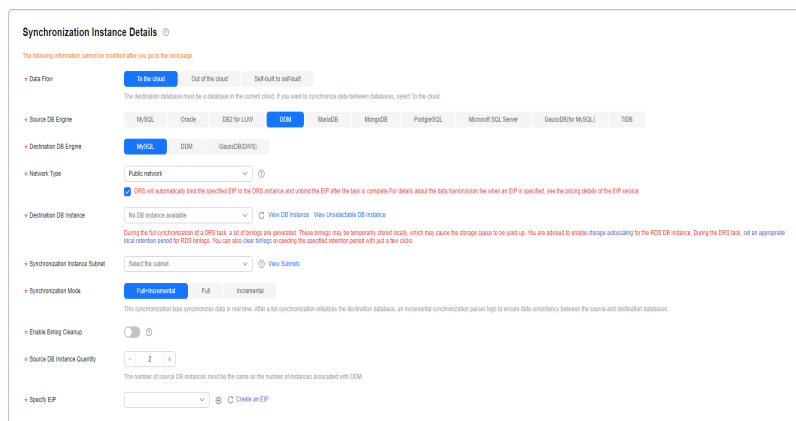


Table 3-232 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select MySQL .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The RDS DB instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Enable Binlog Cleanup	<p>Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.</p>
Source DB Instance Quantity	<p>Specifies the number of DB instances bound to the source DDM database. The default value is 2. The value ranges from 1 to 64. Set this parameter based on the site requirements.</p>

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 3-190 Task type




Table 3-233 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-191 Enterprise Project and Tags




Table 3-234 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-192 Source database information

Source Database

Middleware IP Address or Domain Name

Port

Middleware Username

Database Password

DB Instance	IP Address or Domain Name	Port	Username	Password
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>

This button is available only after the replication instance is created successfully.

Table 3-235 Source database settings

Parameter	Description
Middleware IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Middleware Username	The username of the source DDM instance.
Middleware Password	The password for the source DDM instance username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.
DB Instance	The sharded database details.

 **NOTE**


The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-193 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

Table 3-236 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance you selected when creating the synchronization task. This parameter cannot be changed.

Parameter	Description
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-194 Synchronization Mode

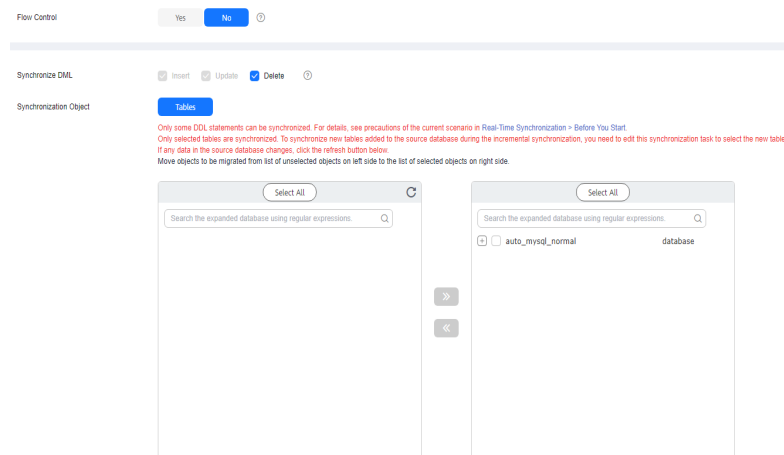
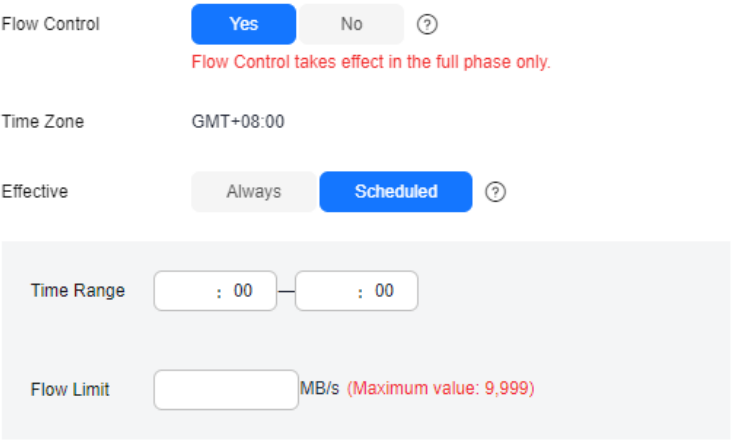



Table 3-237 Synchronization mode and object


Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-195 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 3-196 Data processing



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-197 Task startup settings

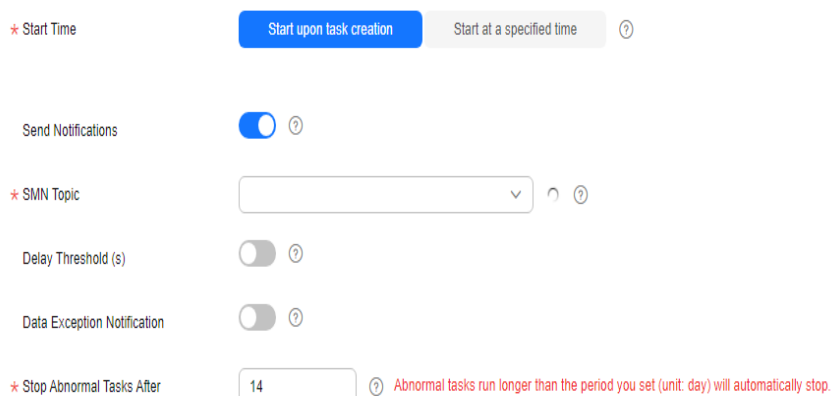



Table 3-238 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.20 From DDM to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-239 Supported databases

Source DB	Destination DB
DDM instances	GaussDB(DWS) clusters (versions 8.1.3 and 8.2.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-240](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-240 Database account permission

Type	Full+Incremental
Source database user	<ul style="list-style-type: none"> • The user of the source DDM database must have at least one permission, for example, SELECT. • The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions.
Destination database user	The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-241](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-241 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Full synchronization supports the synchronization of data, table structures, and common indexes (B-Tree indexes). • The destination database does not support tables with unique keys. During synchronization, the table with unique keys in the source database is ignored. • A table without a primary key cannot be synchronized. If the selected table does not have a primary key, the synchronization fails. • Unique constraints and indexes cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-242](#).

Table 3-242 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> ● The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: '<>\'\" ● Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout. ● The source physical sharded database does not support the enumeration and set types. ● The default value in the timestamp column of the source database must be within the range allowed by the destination database. Otherwise, the synchronization fails.
Use restrictions	<p>General</p> <ul style="list-style-type: none"> ● During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. <p>Full synchronization</p> <ul style="list-style-type: none"> ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. ● If you need to modify the structure of the source table to be synchronized, you must modify the corresponding destination table structure. ● Some DDL statements are supported, including ADD COLUMN, CREATE TABLE, MODIFY COLUMN, CREATE INDEX, DROP INDEX and RENAME INDEX. ● Perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. ● The source database cannot be restored. ● Tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated on the destination database. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If the source database contains a duplicate primary key, data synchronized to the destination database will be less than that in the source database. Therefore, you must check and correct the data before starting the synchronization task. ● If a table in the source database contains system columns _ddm_lock and _slot, data in these columns will be filtered out during full synchronization and incremental synchronization, resulting in data loss. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● In the source database, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. ● When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-198 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

Task Name: [Text input: DRS-5678]

Description: [Text area: 0/256]

Table 3-243 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-199 Synchronization instance information

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

• Data Flow: [To the cloud] | Out of the cloud | **Default to default**

The destination database must be a database in the current cloud. If you want to synchronize data between databases, select the cloud.

• Source DB Engine: MySQL, Oracle, DEZ for LUW, **DCM**, MariaDB, MongoDB, PostgreSQL, Microsoft SQL Server, GaussDB for MySQL, TDengine

• Destination DB Engine: MySQL, EDW, **GaussDB for MySQL**

The syntax of databases such as MySQL and Oracle is different from that of GaussDB(DWS). DDL statements may fail to be synchronized. Contact DWS experts to evaluate DDL support before the synchronization.

• Network Type: Public network

DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance: No DB instance available | [View DB Instance](#) | [View Unavailable DB Instance](#)

• Synchronization Instance Subnet: Select the subnet | [View Subnets](#)

• Synchronization Mode: **Full Synchronization**

This synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Source DB Instance Quantity: 2

The number of source DB instances must be the same as the number of instances associated with DCM.

• Specify EIP: [Dropdown menu] | [Create an EIP](#)

Table 3-244 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB(DWS) instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Quantity	<p>Specifies the number of DB instances bound to the source DDM database. The default value is 2. The value ranges from 1 to 64. Set this parameter based on the site requirements.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-200 Task type




Table 3-245 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-201 Enterprise Project and Tags




Table 3-246 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-202 Source database information

Source Database

Middleware IP Address or Domain Name

Port

Middleware Username

Database Password

DB Instance

IP Address or Domain Name	Port	Username	Password
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>

This button is available only after the replication instance is created successfully.

Table 3-247 Source database settings

Parameter	Description
Middleware IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Middleware Username	The username of the source DDM instance.
Middleware Password	The password for the source DDM instance username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.
DB Instance	The sharded database details.

NOTE

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 3-203 Destination database information

Destination Database

DB Instance Name

Database Username


Database Password 

Table 3-248 Destination database settings

Parameter	Description
DB Instance Name	The default value is the GaussDB(DWS) instance selected for creating the migration task. It cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-204 Synchronization Mode

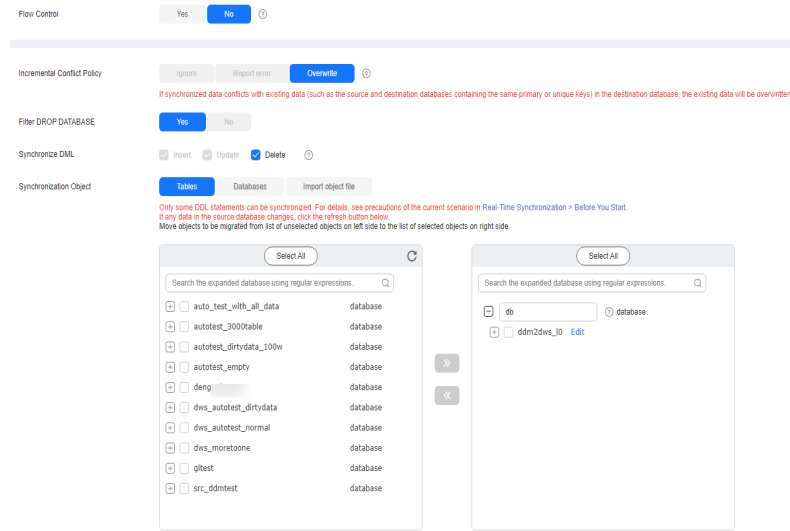
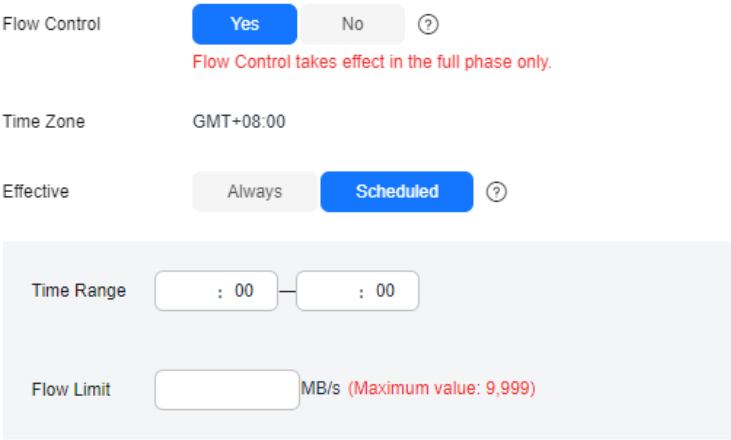



Table 3-249 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-205 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.
Filter DROP DATABASE	During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Synchronization Object	The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables , Import object file , or Databases for Synchronization Object as required. <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Data Processing** page, select the table object to be processed, enter the column name, type, and operation type to be added, confirm the information, and click **Next**. You can configure related rules by referring to [Processing Data](#).

Figure 3-206 Processing data

Additional Columns

You can use additional columns to avoid data conflicts in many-to-one operations.

Belonged Database Table	New name	Column Name	Operation Type	Type	Operation
dds2dms_01_db_full_primary_001_table_01	dds2dms_01_db_full_primary_001_table_01	-	-	-	Add
dds2dms_01_db_full_primary_002_table	dds2dms_01_db_full_primary_002_table	-	-	-	Add
dds2dms_01_db_full_xin_singal_table01	dds2dms_01_db_full_xin_singal_table01	-	-	-	Add
dds2dms_01_db_full_xin_singal_table02	dds2dms_01_db_full_xin_singal_table02	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-207 Task startup settings

* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ?

Delay Threshold (s) ?


Data Exception Notification ?

* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 3-250 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.21 From DDM to DDM

Supported Source and Destination Databases

Table 3-251 Supported databases

Source DB	Destination DB
DDM instances	DDM instances

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-252](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-252 Database account permission

Type	Full+Incremental
Source database user	<ul style="list-style-type: none"> • The user of the source DDM database must have at least one permission, for example, SELECT. • The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions.
Destination database user	The user must have the following permissions: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT. In addition, grant the SELECT permission on all tables. The DDM destination database user must have permissions on the database to be synchronized.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-253](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-253 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Full synchronization supports the synchronization of data, table structures, and indexes.• Only MyISAM and InnoDB tables can be synchronized.• The source database cannot contain tables whose sharding keys are of the timestamp type.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-254](#).

Table 3-254 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: <code>.'<>/\</code>• Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout.

Type	Restrictions
Use restrictions	<p>General</p> <ul style="list-style-type: none"> • During the task startup or full synchronization, you are not advised to perform DDL operations, such as deletion, on the source database. Otherwise, the synchronization may fail. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During synchronization, do not modify the table structure to be synchronized in the source database. • During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • The source database cannot be restored. • Some DDL statements are supported, including ADD COLUMN, CREATE TABLE, CREATE INDEX, DROP INDEX, RENAME INDEX and ADD INDEX. • Perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. ● Create a schema for the destination database in advance. ● Ensure that the destination database is empty before starting the synchronization. Otherwise, data in the destination may be overwritten during incremental synchronization. ● If the destination instance uses columns of the <code>TIMESTAMP</code> or <code>DATETIME</code> data type as its sharding key, the seconds precision of the column is removed after the synchronization. ● The AUTO_INCREMENT value of a table in the destination cannot be less than that of a table in the source. ● During an incremental synchronization of table-level objects, renaming tables is not recommended. ● When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours. ● If the destination DDM version is later than 3.0.4.1, DRS automatically updates the start value of the DDM sequence when the task is complete. ● The destination database user cannot be an administrator.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-208 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page.
 The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-255 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-209 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Task Name: ⓘ

Region:

Project:

Task Name: ⓘ

Description: ⓘ

0/256

Table 3-256 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select DDM .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The DDM instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Quantity	Specifies the number of DB instances bound to the source DDM database. The default value is 2 . The value ranges from 1 to 64 . Set this parameter based on the site requirements.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 3-210 Task type




Table 3-257 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-211 Enterprise Project and Tags




Table 3-258 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-212 Source database information

Source Database

Middleware IP Address or Domain Name

Port

Middleware Username

Database Password

DB Instance

IP Address or Domain Name	Port	Username	Password
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>

This button is available only after the replication instance is created successfully.

Table 3-259 Source database settings

Parameter	Description
Middleware IP Address or Domain Name	The IP address or domain name of the source DDM middleware.
Port	The port of the source DDM middleware. Value range: 1 to 65535
Middleware Username	The username of the source DDM instance.
Middleware Password	The password for the source DDM instance username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.
DB Instance	The sharded database details.

NOTE

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 3-213 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-260 Destination database settings

Parameter	Description
DB Instance Name	The DDM instance you selected when you create the synchronization task. The instance name cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-214 Synchronization Mode

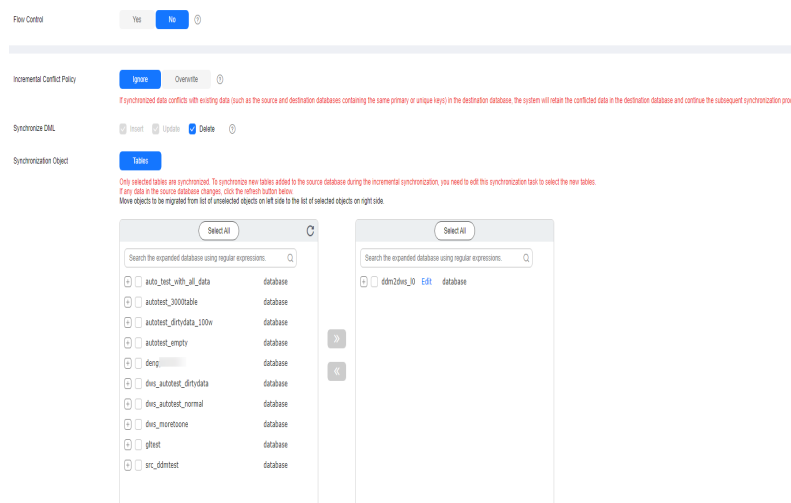
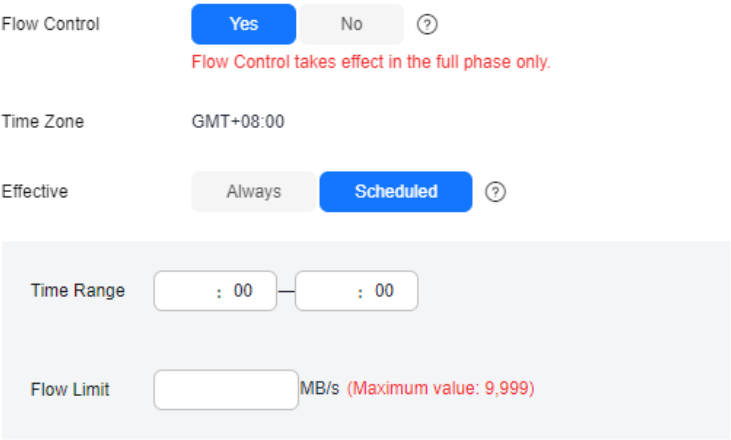



Table 3-261 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-215 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-216 Task startup settings

Table 3-262 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.22 From DB2 for LUW to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-263 Supported databases

Source DB	Destination DB
DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB Primary/Standby (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-264](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-264 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The user must have the CONNECT and DATAACCESS permissions.	The user must have the DBADM permission.
	If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table.	

Type	Full Synchronization	Full+Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-265](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-265 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. ● Tables can be synchronized in real time during incremental synchronization. ● The maximum precision supported by the TIMESTAMP type is 6. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● User-defined types are not supported. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-266](#).

Table 3-266 Precautions

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. The source database cannot contain uncommitted DDL transactions. <p>Full synchronization</p> <ul style="list-style-type: none"> Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. <p>Incremental synchronization</p> <ul style="list-style-type: none"> During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. You can add additional objects in the incremental synchronization phase. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● Arm VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the table structure contains the DB2SECURITYLABEL data, reading data in the full phase may be affected by label-based access control (LBAC) of DB2. As a result, data in the source database is inconsistent with that in the destination database. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● Before starting a full+incremental task, ensure that the source database does not contain uncommitted transactions. You are advised to stop writing data to the source database and then start the task. ● After a task is started, the detach operation on a partition table may cause data inconsistency.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-217 Synchronization task information

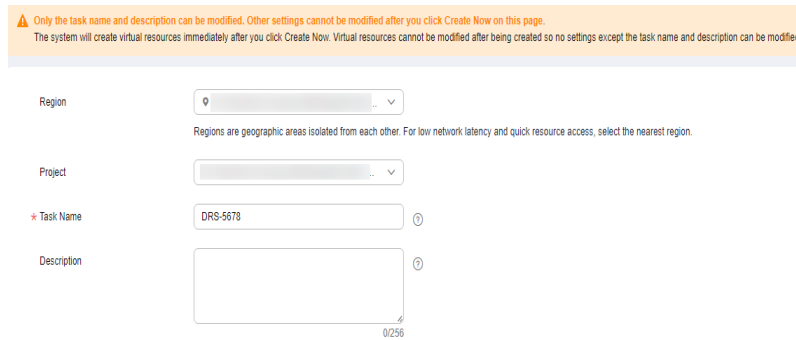


Table 3-267 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-218 Synchronization instance details




Table 3-268 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DB2 for LUW .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles .

Parameter	Description
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-219 AZ




Table 3-269 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-220 Enterprise Project and Tags




Table 3-270 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 3-221 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized--

This button is available only after the replication instance is created successfully.

Table 3-271 Source database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the source DB2 for LUW database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source DB2 for LUW database.
Database Password	The password for the source database username.
Database Name	The name of the database to which the synchronization objects belong in the source DB2 for LUW.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> • The maximum size of a single certificate file that can be uploaded is 500 KB. • If the SSL certificate is not used, your data may be at risk. • Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth • DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures.
Synchronize Driver	You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection.

 **NOTE**

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-222 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 3-272 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for logging in to the destination database.
Database Password	Password for the destination database username. The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 3-223 Synchronization mode

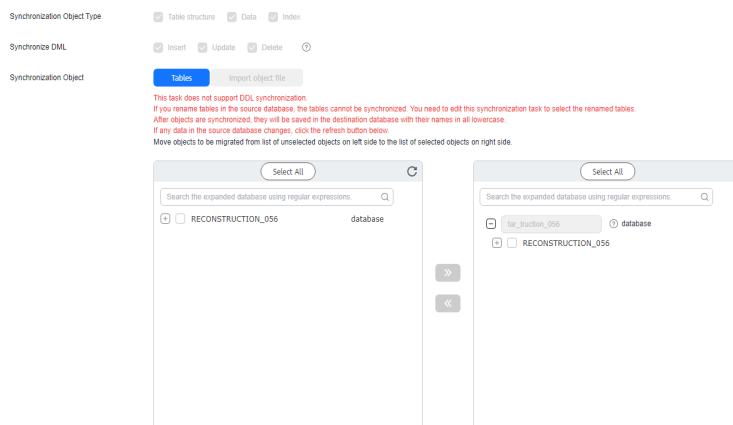



Table 3-273 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can synchronize tables or import object files based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-224 Task startup settings





Table 3-274 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.23 From DB2 for LUW to GaussDB Distributed

Supported Source and Destination Databases

Table 3-275 Supported databases

Source DB	Destination DB
DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB Distributed (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-276](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-276 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The user must have the CONNECT and DATAACCESS permissions.	The user must have the DBADM permission.
	If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table.	

Type	Full Synchronization	Full+Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-277](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-277 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized.• Tables can be synchronized in real time during incremental synchronization.• The maximum precision supported by the TIMESTAMP type is 6.• The LOB type supported by incremental synchronization cannot exceed 10 MB.• User-defined types are not supported.• During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes.• VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization.• Tables whose default values contain expression functions cannot be synchronized.• Temporary tables in the source database cannot be synchronized.• DDL statements executed in the source database cannot be synchronized.• For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized.• If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-278](#).

Table 3-278 Precautions

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. The source database cannot contain uncommitted DDL transactions. <p>Full synchronization</p> <ul style="list-style-type: none"> Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. <p>Incremental synchronization</p> <ul style="list-style-type: none"> During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. You can add additional objects in the incremental synchronization phase. During an incremental synchronization, if you update the primary key column or the first column of a table that does not have a primary key in source database DB2 for LUW, the distribution column will be updated in the GaussDB database, which may cause data inconsistency. Thus, do not update the preceding columns. <p>Troubleshooting</p>

Type	Restrictions
	<ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.
Other restrictions	<ul style="list-style-type: none"> ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● Arm VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the table structure contains the DB2SECURITYLABEL data, reading data in the full phase may be affected by label-based access control (LBAC) of DB2. As a result, data in the source database is inconsistent with that in the destination database. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● After a task is started, the detach operation on a partition table may cause data inconsistency.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-225 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region ..
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project ..

* Task Name ⓘ

Description ⓘ
0/256

Table 3-279 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-226 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud Self-built to self-built

The destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

• Source DB Engine Oracle GaussDB MySQL MariaDB PostgreSQL GaussDB(MySQL)

• Destination DB Engine GaussDB Oracle

• Network Type ⓘ

DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance ⓘ View DB Instance View Unselected DB Instance

• Synchronization Instance Subnet ⓘ View Subnets View Occupied IP Address

• Synchronization Mode Full-replica Full

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization takes logs to ensure data consistency between the source and destination databases.

• Specify EIP ⓘ Create an EIP

Table 3-280 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DB2 for LUW .
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles .

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 3-227 AZ




Table 3-281 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-228 Enterprise Project and Tags




Table 3-282 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 3-229 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized--

This button is available only after the replication instance is created successfully.

Table 3-283 Source database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the source DB2 for LUW database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source DB2 for LUW database.
Database Password	The password for the source database username.
Database Name	The name of the database to which the synchronization objects belong in the source DB2 for LUW.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none">• The maximum size of a single certificate file that can be uploaded is 500 KB.• If the SSL certificate is not used, your data may be at risk.• Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth• DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures.
Synchronize Driver	You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection.

 NOTE

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-230 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-284 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for logging in to the destination database.
Database Password	Password for the destination database username. The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 3-231 Synchronization mode

Synchronization Object Type Table structure Data Index

Synchronize DML Insert Update Delete

Synchronization Object Tables

This task does not support DDL synchronization. If you rename tables in the source database, the tables cannot be synchronized. You need to edit this synchronization task to select the renamed tables. After objects are synchronized, they will be saved in the destination database with their names in all lowercase. If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

Search the expanded database using regular expressions:

RECONSTRUCTION_056 database


Select All

Search the expanded database using regular expressions:

tar_function_056 database

RECONSTRUCTION_056

Table 3-285 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-232 Task startup settings

Table 3-286 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.24 From DB2 for LUW to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-287 Supported databases

Source DB	Destination DB
DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB(DWS) (versions 8.1.3 and 8.2.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-288](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-288 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The user must have the CONNECT and DATAACCESS permissions.	The user must have the DBADM permission.
	If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table.	
Destination database user	The user must have the following permissions: <ul style="list-style-type: none">• If the destination side does not contain databases, schemas, or tables, the destination database user must have the permission to create databases, the permission to create schemas in a database, or the permission to create tables in a schema.• The INSERT, SELECT, UPDATE, and DELETE permissions are required for each table.	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-289](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-289 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. ● Tables can be synchronized in real time during incremental synchronization. ● The maximum precision supported by the TIMESTAMP type is 6. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● User-defined types are not supported. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. ● The number of synchronization objects of a single schema in the source database cannot exceed 32,766. ● If the index name in the source database contains more than 63 characters, the index will be automatically renamed when it is being replayed in the destination database. ● In incremental synchronization, the Datacapture attribute of the table must be set to Y.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-290](#).

Table 3-290 Precautions

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. The source database cannot contain uncommitted DDL transactions. <p>Full synchronization</p> <ul style="list-style-type: none"> Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During table structure synchronization in the full phase, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. You can add additional objects in the incremental synchronization phase. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. ● Arm VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the table structure contains the DB2SECURITYLABEL data, reading data in the full phase may be affected by label-based access control (LBAC) of DB2. As a result, data in the source database is inconsistent with that in the destination database. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● The DECFLOAT data on the source database is a decimal floating point number. For DECFLOAT(16) data, the minimum exponent is -383 and the maximum exponent is 384. For DECFLOAT(34) data, the minimum exponent is -6143 and the maximum exponent is 6144. When DRS is used to migrate table structures, the DECFLOAT type will be converted to the NUMERIC (65,10) type of GaussDB(DWS). During data synchronization, an error may be reported indicating that data cannot be written due to low precision. In this case, you

Type	Restrictions
	<p>need to manually change the precision of the corresponding column in the destination database.</p> <ul style="list-style-type: none"> • Before starting a full+incremental task, ensure that the source database dose not contain uncommitted transactions. You are advised to stop writing data to the source database and then start the task. • After a task is started, the detach operation on a partition table may cause data inconsistency.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-233 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

Task Name ⓘ

Description ⓘ
0/256

Table 3-291 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-234 Synchronization instance information

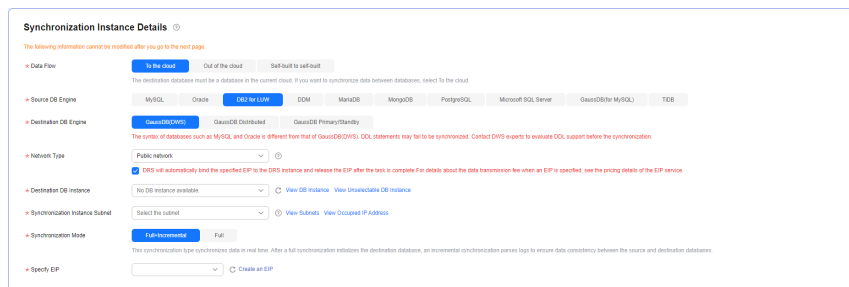


Table 3-292 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is Databases on the current cloud .
Source DB Engine	Select DB2 for LUW .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The GaussDB(DWS) instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-235 AZ

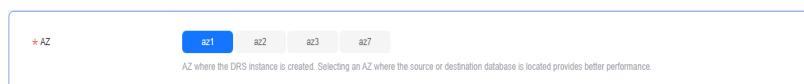


Table 3-293 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-236 Enterprise Project and Tags




Table 3-294 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 3-237 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronize Driver Synchronized -

Test Connection This button is available only after the replication instance is created successfully.

Table 3-295 Source database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the source DB2 for LUW database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source DB2 for LUW database.
Database Password	The password for the source database username.
Database Name	The name of the database to which the synchronization objects belong in the source DB2 for LUW.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures.

Parameter	Description
Synchronize Driver	You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection.

 **NOTE**


The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-238 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 


 Test successful

Table 3-296 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(DWS) instance that you selected during task creation and cannot be changed.
Database Username	The username for logging in to the destination database.
Database Password	Password for the destination database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 3-239 Synchronization Mode

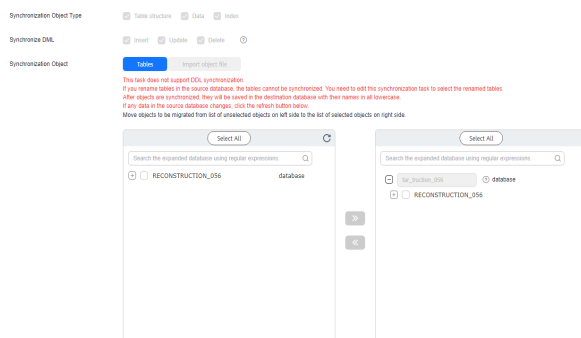



Table 3-297 Synchronization object

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

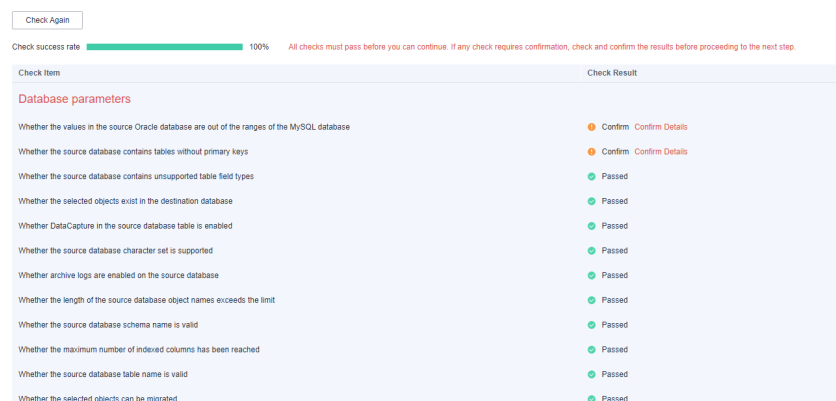
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 3-240 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-241 Task startup settings

Table 3-298 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.25 From TiDB to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-299 Supported databases

Source DB	Destination DB
TiDB 4.0.0 and later (excluding the development version)	GaussDB(for MySQL) Primary/Standby (version 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-300](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-300 Database account permission

Type	Full+Incremental
Source database user	SELECT and CONFIG permissions.
Destination database user	SELECT, CREATE, DROP, INSERT, DELETE, UPDATE, ALTER, REFERENCES and INDEX permissions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-301](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-301 Supported synchronization objects

Type	Precautions
Synchronization objects	Table structures, data, and indexes and constraints of selected tables can be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-302](#).

Table 3-302 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The names of the source objects to be synchronized cannot contain non-ASCII characters or the following special characters: '<'>\/ • The Pump and Drainer components must be deployed on the source database. For details, see TiDB Binlog Cluster Deployment. • The output type of the Drainer component must be set to Kafka. For details, see Configuring Kafka Drainer • The precision value range of the double type of TiDB is different from that of GaussDB(for MySQL). Tables with primary keys in the source database cannot use columns of the double type as primary key columns, and tables without primary keys in the source database cannot contain columns of the double type. If data of the double type is contained in the preceding scenarios, data in the source database may be inconsistent with that in the destination database. <p>CAUTION</p> <ul style="list-style-type: none"> • TiDB Binlog is incompatible with the following features in TiDB 5.0: <ul style="list-style-type: none"> TiDB clustered index: After TiDB Binlog is enabled, you cannot create a clustered index that is defined as a single non-integer primary key. The insert, delete, and update operations on existing tables with clustered indexes are not synchronized to the downstream applications through TiDB Binlog. To synchronize a table with clustered indexes, upgrade the TiDB version to 5.1. TiDB system variable tidb_enable_async_commit: After TiDB Binlog is enabled, the performance cannot be improved if this option is enabled. TiDB system variable tidb_enable_1pc: After TiDB Binlog is enabled, the performance cannot be improved if this option is enabled. • TiDB Binlog is incompatible with the following features in TiDB 4.0.7: <ul style="list-style-type: none"> TiDB system variable tidb_enable_amend_pessimistic_txn: After TiDB Binlog is enabled, the data replicated from the TiDB Binlog may be inconsistent.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <p>Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. Some DDL statements, including CREATE TABLE, DROP TABLE, ALTER TABLE, RENAME TABLE and TRUNCATE TABLE, can be synchronized. <p>Synchronization comparison</p> <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If lower_case_table_names of the destination database is set to 1, databases or tables containing uppercase letters cannot be synchronized. ● In full+incremental synchronization, if the destination database already has a table structure, ensure that the table structure is the same as that on the source database or contains columns in the source database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● Triggers enabled in the destination database cannot be associated with synchronization tables. ● You can add additional objects in the incremental synchronization phase. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail.

Procedure

This section uses TiDB to GaussDB(for MySQL) synchronization to the cloud as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-242 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-5678

Description: [Text area]

0/256

Table 3-303 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-243 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: **To the cloud** (Selected) | Out of the cloud | Self-built to self-built

Source DB Engine: **TiDB** (Selected) | MySQL | Oracle | DB2 for LUW | DDM | MariaDB | MongoDB | PostgreSQL | Microsoft SQL Server | GaussDB for MySQL

Destination DB Engine: **GaussDB for MySQL** (Selected)

Network Type: **PUBLIC network** (Selected) | Private network

Destination DB Instance: **No DB instance available** (Selected) | View DB Instance | View Unavailable DB Instance

Synchronization Instance Subnet: **Select the subnet** (Selected) | View Subnets | View Occupied IP Address

Synchronization Mode: **Full incremental** (Selected)

Spiky EP: [Dropdown menu] | Create an EP

Table 3-304 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select TiDB .

Parameter	Description
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	<p>Available options: Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The GaussDB(for MySQL) instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-244 AZ




Table 3-305 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-245 Enterprise Project and Tags




Table 3-306 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify Kafka, source and destination database information and click **Test Connection** for Kafka and the source and destination databases to check whether they have been connected. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 3-246 Source database information


Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection


 Test successful

Table 3-307 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source TiDB database.
Port	The port number must be an integer ranging from 1 to 65535, and the default value is 4000.
Database Username	The username for accessing the source TiDB database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Figure 3-247 Kafka information



Table 3-308 Kafka information

Parameter	Description
IP Address or Domain Name	IP address or domain name of Kafka DRS obtains the binlog generated by TiDB from Kafka and synchronizes the incremental data to the destination database in real time.

Figure 3-248 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-309 Destination database settings

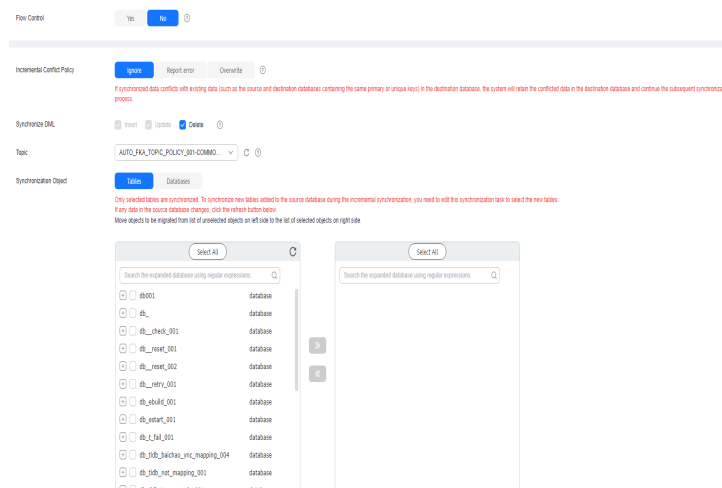
Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, topic, and synchronization object, and click **Next**.

Figure 3-249 Synchronization Mode



Flow Control Yes No

Incremental Conflict Policy Ignore Report error Overwrite

Synchronization DML Insert Update Delete

Topic AUTO_FKA_TOPIC_POLICY_001-COMMO

Synchronization Object Tables Databases

Only selected tables are synchronized. To synchronize new tables in the source database during the incremental synchronization, you need to add the synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. Please objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

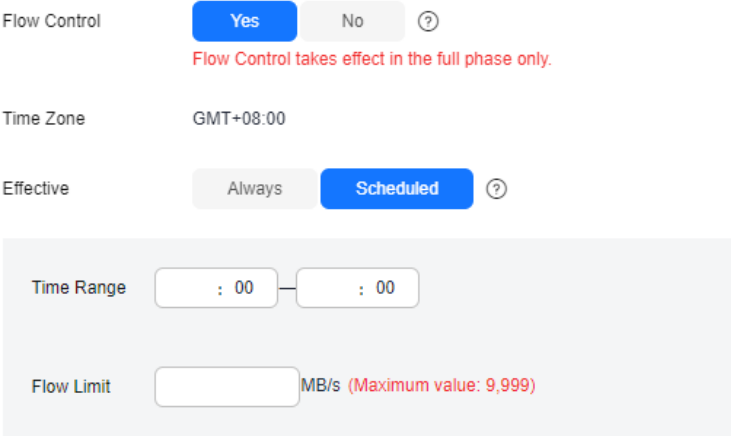
Search the expanded database using regular expressions


- db001 database
- db_ database
- db_466_001 database
- db_invest_001 database
- db_invest_002 database
- db_refu_001 database
- db_sdm_001 database
- db_sdm_001 database
- db_sdm_360cha_001_mapping_004 database
- db_sdm_001_mapping_001 database
- db_sdm_new_consume_001 database

Select All

Search the expanded database using regular expressions

Table 3-310 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-250 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Topic	<p>Select a topic that stores TiDB binlogs in the Kafka, or synchronization will fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-251 Task startup settings

Table 3-311 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.26 From Microsoft SQL Server to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-312 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	GaussDB(DWS) clusters (versions 8.1.3 and 8.2.0)

 **NOTE**

Only whitelisted users can use this function.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-313](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-313 Database account permission

Type	Full+Incremental Synchronization
Source database user	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	<ul style="list-style-type: none"> • If the destination instance does not contain the database to be synchronized, the CREATEDB permission is required. • If the destination instance contains databases but does not have schemas, the CONNECT and CREATE permissions for the databases are required. • If the destination instance contains databases and schemas but does not contain tables, the CONNECT permission for databases and the USAGE and CREATE permissions for schemas are required. • If the destination instance contains databases, schemas, and tables, the CONNECT permission for databases, the USAGE permission for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, REFERENCES, and INDEX permissions for tables are required.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-314](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-314 Supported synchronization objects

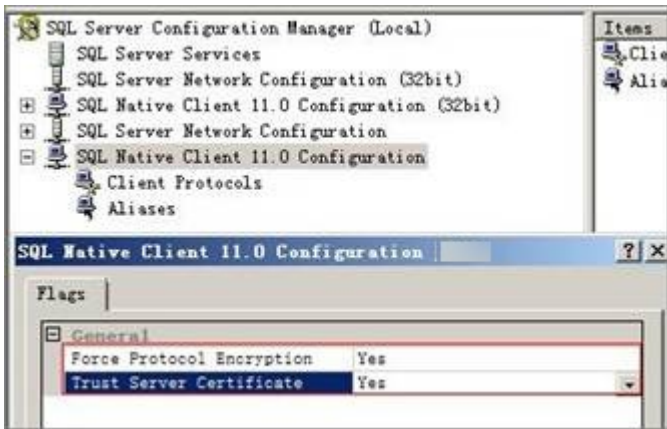
Type	Precautions
Objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized.

Type	Precautions
	<p>Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail.</p> <ul style="list-style-type: none"> - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. • Scope of incremental synchronization <ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported - DDL statements are not supported. - In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-315](#).

Table 3-315 Precautions

Type	Restrictions
<p>Restrictions on the source database</p>	<p>If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-252.</p> <p>Figure 3-252 Client configuration</p> 

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <p>Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● DML statements, including INSERT, UPDATE, and DELETE, are supported. ● DDL operations performed on the source database will not be synchronized to the destination database. ● The IMAGE, TEXT, and NTEXT big data types cannot be deleted. ● You can add additional synchronization objects. <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> – The default value may be left empty. – The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. ● The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. ● Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks. ● B-Tree in GaussDB(DWS) can create indexes on a maximum of 32 fields. If the number of fields on which combined indexes are created in the source database exceeds 32, indexes fails to be created in the destination database.

Procedure

This section uses Microsoft SQL Server to GaussDB(DWS) synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-253 Synchronization task information

⚠️ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] .. v
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu] .. v

* Task Name: DRS-5678 ⓘ

Description: [Text area] ⓘ
0/256

Table 3-316 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-254 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** | Out of the cloud | Self-built to self-built

The destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

• Source DB Engine: MySQL | Oracle | DB2 for LUW | DDM | MaxDB | KingDB | PostgreSQL | **Microsoft SQL Server** | GaussDB for MySQL | TCR

• Destination DB Engine: **GaussDB OMS** | GaussDB Distributed | GaussDB Primary/Standby

The syntax of databases such as MySQL and Oracle is different from that of GaussDB OMS. DDL statements may fail to be synchronized. Contact DRS experts to evaluate DDL support before the synchronization.

• Network Type: Public network ⓘ

DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance: No DB instance available | View DB Instance | View Unavailable DB Instance

• Synchronization Instance Subnet: Select the subnet | View Subnets | View Occupied IP Address

• Synchronization Mode: **Full incremental**

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Specify EIP: [Dropdown menu] | Create an EIP

Table 3-317 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The destination is a GaussDB(DWS) instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 3-255 AZ

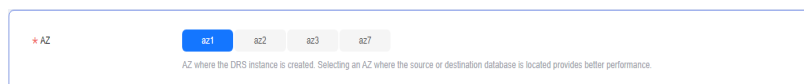


Table 3-318 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-256 Enterprise Project and Tags




Table 3-319 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 3-257 Source database information

Source Database

DRS migrates only some key parameters to the destination database. For the other parameters that cannot be migrated, you need to use [parameter templates](#) to configure them on the destination database.

Database Type: self-built database RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-320 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

Figure 3-258 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

Test successful

Table 3-321 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(DWS) instance that you selected during task creation and cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-259 Synchronization mode

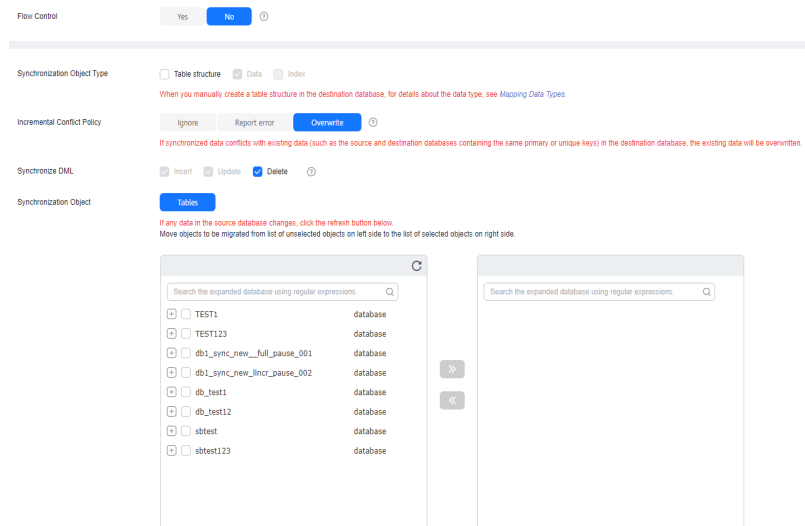
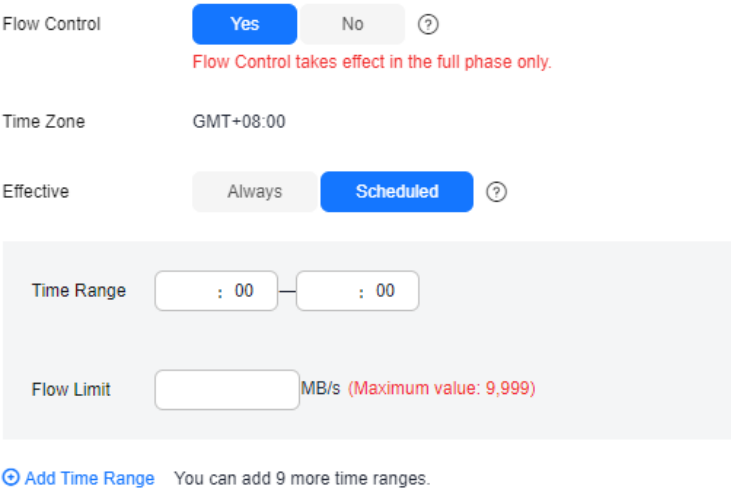



Table 3-322 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-260 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-261 Task startup settings

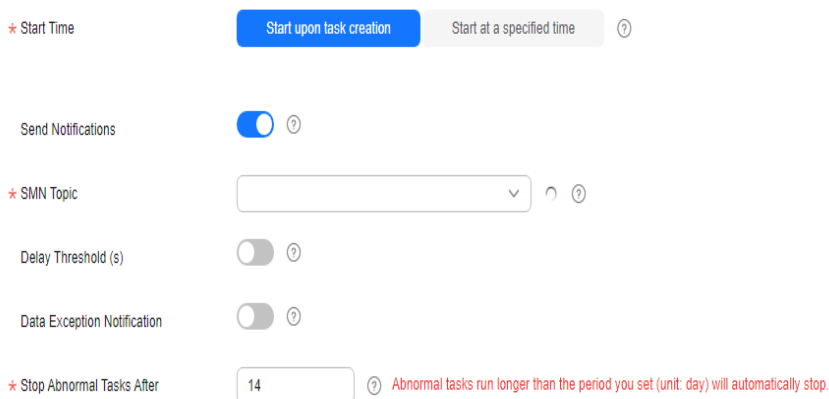



Table 3-323 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.27 From Microsoft SQL Server to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-324 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022)ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022)Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022)RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022)	GaussDB Primary/Standby (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-325](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-325 Database account permission

Type	Full	Incremental and Full+Incremental
Source database user	Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	<ul style="list-style-type: none"> Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

Supported Synchronization Objects

Table 3-326 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-326 Supported synchronization objects

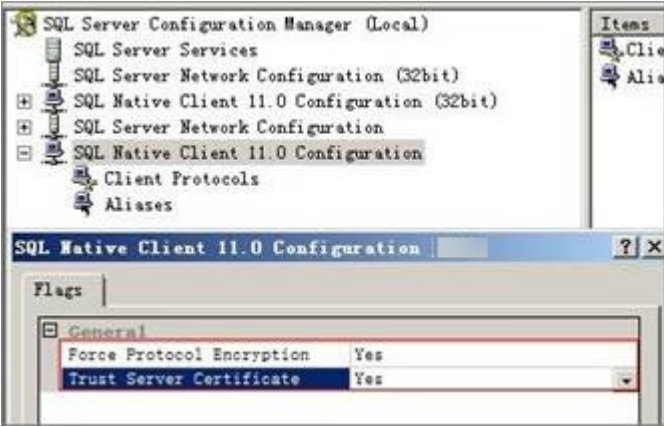
Type	Precautions
Objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized. - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported. - In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-327](#).

Table 3-327 Precautions

Type	Restrictions
Restrictions on the source database	<p>If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-262.</p> <p>Figure 3-262 Client configuration</p> 

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <p>Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> DML statements, including INSERT, UPDATE, and DELETE, are supported. DDL operations performed on the source database will not be synchronized to the destination database. The IMAGE, TEXT, and NTEXT big data types cannot be deleted. Tables whose primary key type is BLOB, TEXT, CLOB, NCLOB, or BYTEA cannot be deleted or updated. When adding synchronization objects in an incremental task, you are not advised to perform DML operations on the newly added synchronization tables before the subtask change is complete. Otherwise, some data may fail to be synchronized to the destination database. <p>Synchronization comparison</p> <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. ● The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. ● Do not use the date type as a primary key in a database. Due to precision loss, this feature affects the synchronization of UPDATE and DELETE statements in DRS incremental synchronization. ● If the destination database is in Oracle compatibility mode, empty strings written to the destination database are processed as NULL. If the source database contains empty strings and has the NOT NULL constraint, data fails to be written to the destination database. ● During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. ● Object names will be converted to lowercase letters after being synchronized to the destination database. Therefore, the selected source database tables cannot contain tables with the same name but different letter cases. Otherwise, the synchronization fails. ● Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.

Procedure

This section uses Microsoft SQL Server to GaussDB primary/standby synchronization to the cloud as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-263 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page.
 The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-328 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-264 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow
 To the cloud
 Out of the cloud
 Self-built to self-built

The destination databases must be a database in the current cloud. If you want to synchronize data between databases, select to the cloud.

• Source DB Engine
 MySQL
 Oracle
 DB2 for LUW
 EDM
 MaxDB
 MongoDB
 PostgreSQL
 Microsoft SQL Server
 GaussDB for MySQL
 TDB

• Destination DB Engine
 GaussDB(DWS)
 GaussDB Distributed
 GaussDB Primary/Standby

• Network Type
 ⓘ

DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance
 ⓘ
 [View DB Instance](#)
 [View Unavailable DB Instance](#)

• Synchronization Instance Subnet
 ⓘ
 [View Subnets](#)
 [View Occupied IP Address](#)

• Synchronization Mode
 Full/Incremental
 Full
 Incremental

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Specify EIP
 ⓘ
 [Create an EIP](#)

Table 3-329 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example. <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB primary/standby instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-265 AZ




Table 3-330 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-266 Enterprise Project and Tags




Table 3-331 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 3-267 Source database information

Source Database

DRS migrates only some key parameters to the destination database. For the other parameters that cannot be migrated, you need to use parameter templates to configure them on the destination database.

Database Type: self-built database RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-332 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

Figure 3-268 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

Table 3-333 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-269 Synchronization mode

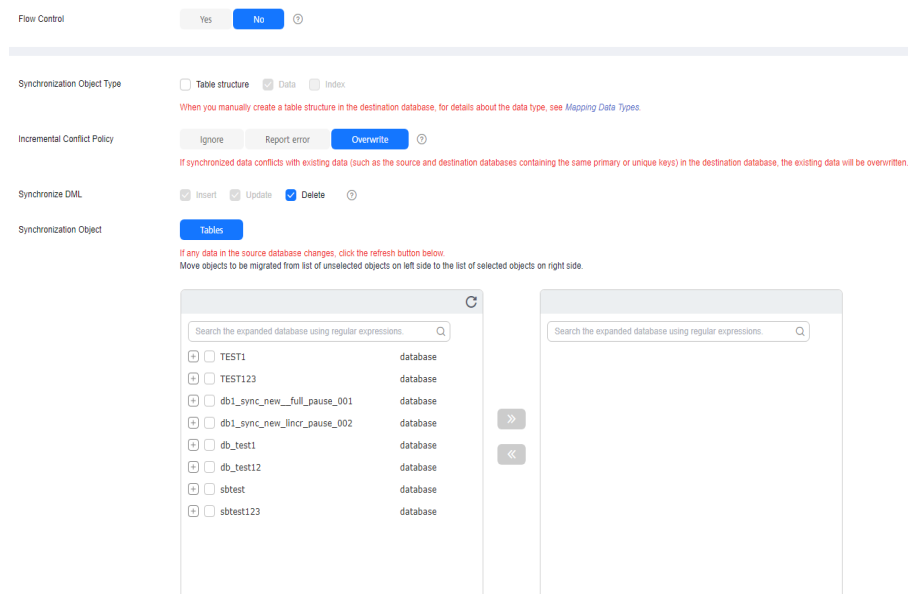
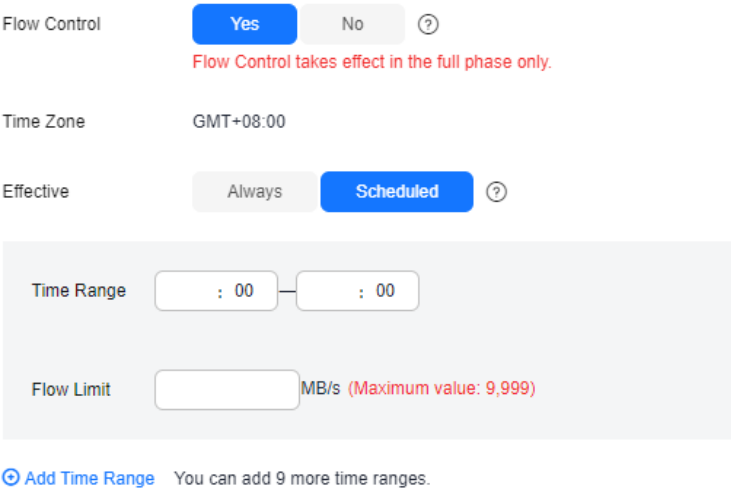



Table 3-334 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-270 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.
Start Point	<p>This option is available if you select Incremental in 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>If not specified, the current LSN of the database is used. For SQL Server, the value is the maximum log sequence number in the start_lsn column in the cdc.lsn_time_mapping system catalog.</p> <p>If specified, the valid LSN range is [minimum LSN reserved for CDC for the table to be synchronized, current LSN of the database]. Data changes after the LSN is specified are synchronized.</p> <p>1. Run the following command to query the minimum LSN reserved for CDC for the table to be synchronized:</p> <pre>select min(start_lsn) as minLsn from cdc.change_tables where source_object_id in (OBJECT_ID('[schemaXXX].[tableXXX]'));</pre> <p>If the data capture instance is not found or the caller is not authorized to access the changed data associated with the data capture instance, 0x000000000000000000000000 is returned.</p> <p>2. Run the following command to query the current LSN of the database:</p> <pre>select sys.fn_cdc_get_max_lsn ();</pre> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-271 Task startup settings

Table 3-335 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.

Parameter	Description
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.28 From Microsoft SQL Server to GaussDB Distributed

Supported Source and Destination Databases

Table 3-336 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	GaussDB Distributed (1.0.0 and later versions)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-337](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-337 Database account permission

Type	Full	Incremental and Full+Incremental
Source database user	Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	<ul style="list-style-type: none"> • Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. • Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. • Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

Supported Synchronization Objects

[Table 3-338](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-338 Supported synchronization objects

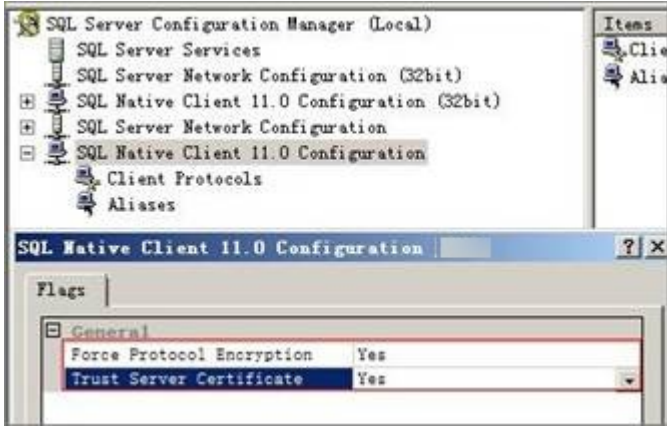
Type	Precautions
Objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized. - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported. - In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-339](#).

Table 3-339 Precautions

Type	Restrictions
Restrictions on the source database	<p>If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-272.</p> <p>Figure 3-272 Client configuration</p> 

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <p>Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● DML statements, including INSERT, UPDATE, and DELETE, are supported. ● DDL operations performed on the source database will not be synchronized to the destination database. ● The IMAGE, TEXT, and NTEXT big data types cannot be deleted. ● Tables whose primary key type is BLOB, TEXT, CLOB, NCLOB, or BYTEA cannot be deleted or updated. ● When adding synchronization objects in an incremental task, you are not advised to perform DML operations on the newly added synchronization tables before the subtask change is complete. Otherwise, some data may fail to be synchronized to the destination database. <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. ● The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. ● Do not use the date type as a primary key in a database. Due to precision loss, this feature affects the synchronization of UPDATE and DELETE statements in DRS incremental synchronization. ● If the destination database is in Oracle compatibility mode, empty strings written to the destination database are processed as NULL. If the source database contains empty strings and has the NOT NULL constraint, data fails to be written to the destination database. ● During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. ● Object names will be converted to lowercase letters after being synchronized to the destination database. Therefore, the selected source database tables cannot contain tables with the same name but different letter cases. Otherwise, the synchronization fails. ● Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.

Procedure

This section uses Microsoft SQL Server to GaussDB distributed synchronization to the cloud as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-273 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 3-340 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-274 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you click the next page.

- Data Flow: To the cloud Out of the cloud Set built to set built
- Source DB Engine: MySQL Oracle DB2 for LUW EDW MaxDB MongoDB PostgreSQL **Microsoft SQL Server** GaussDB for MySQL TCE
- Destination DB Engine: GaussDB(DWS) **GaussDB Distributed** GaussDB Primary/Standby
- Network Type: Public network Private network
 - DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.
- Destination DB Instance: No DB instance available [View DB Instance](#) [View Unavailable DB Instance](#)
- Synchronization Instance Subnet: Select the subnet [View Subnets](#) [View Occupied IP Address](#)
- Synchronization Mode: **Full-incremental** Full Incremental
- Specify EIP: Create an EIP

The synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronizer parses logs to ensure data consistency between the source and destination databases.

Table 3-341 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	An available GaussDB distributed instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-275 AZ




Table 3-342 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-276 Enterprise Project and Tags




Table 3-343 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 3-277 Source database information

Source Database

Table 3-344 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

Figure 3-278 Destination database information

Destination Database

Table 3-345 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-279 Synchronization mode

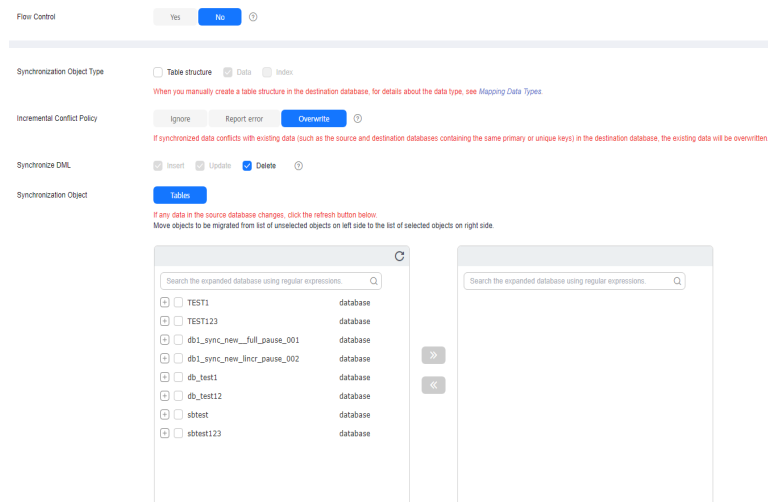
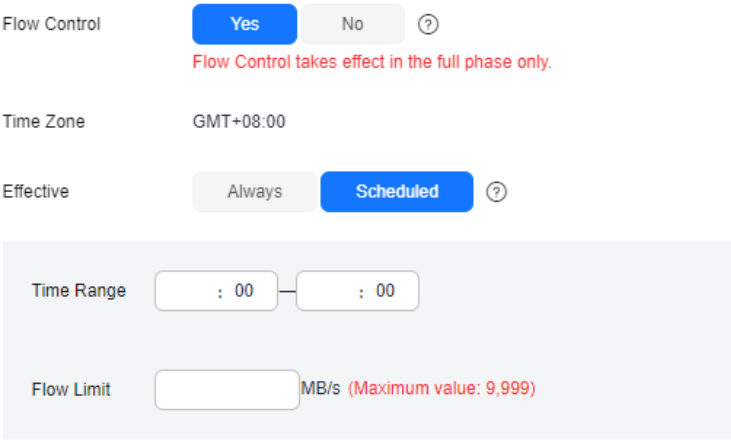



Table 3-346 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-280 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.
Start Point	<p>This option is available if you select Incremental in 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>If not specified, the current LSN of the database is used. For SQL Server, the value is the maximum log sequence number in the start_lsn column in the cdc.lsn_time_mapping system catalog.</p> <p>If specified, the valid LSN range is [minimum LSN reserved for CDC for the table to be synchronized, current LSN of the database]. Data changes after the LSN is specified are synchronized.</p> <p>1. Run the following command to query the minimum LSN reserved for CDC for the table to be synchronized:</p> <pre>select min(start_lsn) as minLsn from cdc.change_tables where source_object_id in (OBJECT_ID('[schemaXXX].[tableXXX]'));</pre> <p>If the data capture instance is not found or the caller is not authorized to access the changed data associated with the data capture instance, 0x000000000000000000000000 is returned.</p> <p>2. Run the following command to query the current LSN of the database:</p> <pre>select sys.fn_cdc_get_max_lsn ();</pre> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-281 Task startup settings

Table 3-347 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.

Parameter	Description
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.29 From Microsoft SQL Server to Microsoft SQL Server

Supported Source and Destination Databases

Table 3-348 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	<p>RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022)</p> <p>NOTE The major version of the destination database must be the same as or later than that of the source database.</p>

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-349](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-349 Database account permission

Type	Full+Incremental
Source database user	At least the sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	<ul style="list-style-type: none"> • If the destination end does not contain databases, the destination database user must have the create any database permission. • If the destination end contains databases, the destination database user must have the connect, create table, alter any schema, and select permissions for the databases.

Supported Synchronization Objects

[Table 3-350](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-350 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ol style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - If the source database contains auto-increment columns and the table structure is selected to be synchronized during a full synchronization, data will be abnormal after auto-increment columns are synchronized to the destination database (auto-increment columns cannot be inserted to the destination database). If the table structure is not selected during a full synchronization and common columns are created in the destination database table structure, data synchronization is normal.

Type	Precautions
	<ul style="list-style-type: none"> - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. • Scope of incremental synchronization <ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported. - In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-351](#).

Table 3-351 Precautions

Type	Restrictions
Restrictions on the source database	<p>If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-282.</p> <p>Figure 3-282 Client configuration</p> 

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. <p>Full synchronization</p> <p>Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> DML statements, including INSERT, UPDATE, and DELETE, are supported. DDL operations performed on the source database will not be synchronized to the destination database. The IMAGE, TEXT, and NTEXT big data types cannot be deleted. <p>Synchronization comparison</p> <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. ● The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. ● Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.

Procedure

This section uses Microsoft SQL Server to Microsoft SQL Server synchronization to the cloud as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-283 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-352 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-284 Synchronization instance details

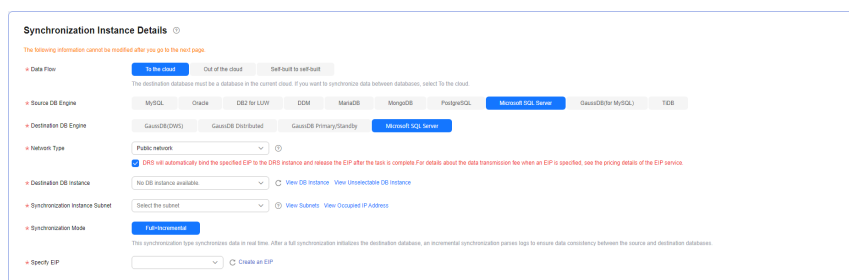


Table 3-353 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select Microsoft SQL Server .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select an RDS for SQL Server instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-285 AZ

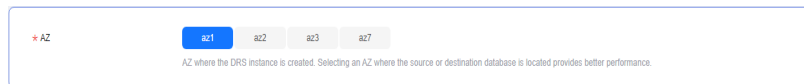


Table 3-354 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-286 Enterprise Project and Tags




Table 3-355 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 3-287 Source database information

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-356 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

Figure 3-288 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-357 Destination database settings

Parameter	Description
DB Instance Name	The RDS for SQL Server instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-289 Synchronization Mode

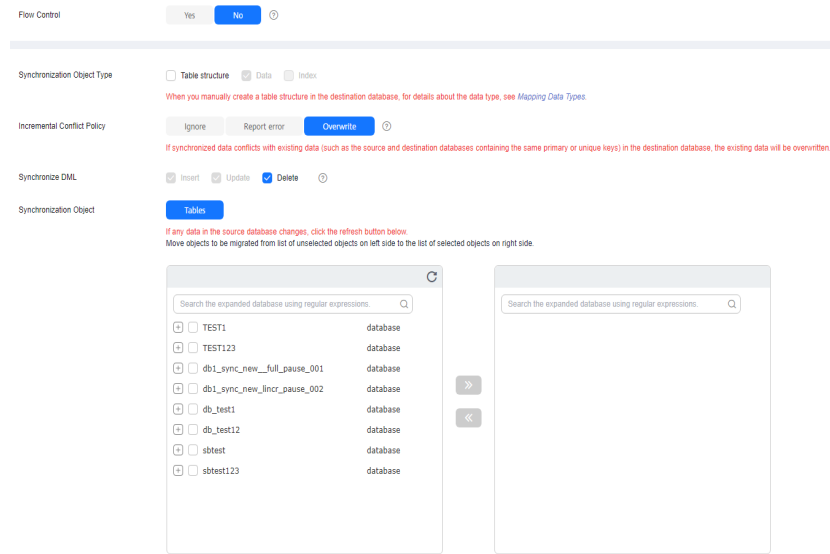
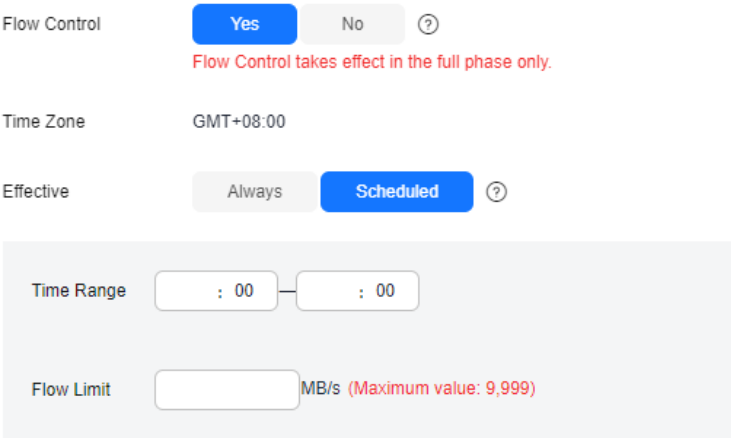



Table 3-358 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-290 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Common index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is mandatory. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-291 Task startup settings





Table 3-359 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.30 From Microsoft SQL Server to MySQL

Supported Source and Destination Databases

Table 3-360 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	RDS for MySQL 5.5, 5.6, 5.7, and 8.0

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-361 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-361 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: HIERARCHYID, TABLE, SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized. - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none">- DML statements, including INSERT, UPDATE, and DELETE, are supported.- DDL statements are not supported.- In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-362](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

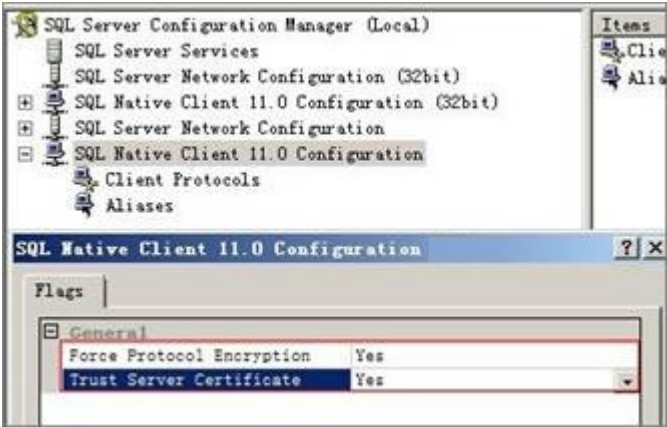
Table 3-362 Database account permission

Type	Full+Incremental Synchronization
Source database user	At least the sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	The user must have the SELECT, CREATE, DROP, INSERT, DELETE, UPDATE, ALTER, REFERENCES and INDEX permissions.

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-363 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - SQL Server Agent must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-292. <p style="text-align: center;">Figure 3-292 Client configuration</p>  ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: [] ? , - The SQL Server character type uses the number of bytes as the length, and the MySQL character type uses the number of characters as the length and the character set is utf8mb4. Therefore, the number of bytes occupied by a field in the destination database is n x 4. - The length of each index cannot exceed 3,072 bytes after being converted into a field in the destination database. - The total number of indexes cannot exceed 64. - The total length of fields after being mapping to the destination database cannot exceed 66,535 bytes. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination DB instance must have sufficient storage space.

Type	Constraints
	<ul style="list-style-type: none"> - The destination DB instance cannot contain databases with the same name as the source database (except system databases). • Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. - Do not use an imprecise value, such as floating-point type and date type, as a primary key or an equivalent condition after WHERE in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data type mapping, see Mapping Data Types from Microsoft SQL Server to MySQL. - Object names will be converted to lowercase letters after being synchronized to the destination database. Therefore, the selected source database tables cannot contain tables with the same name but different letter cases. Otherwise, the synchronization fails. - Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. DDL operations performed on the source database will not be synchronized to the destination database. The IMAGE, TEXT, and NTEXT big data types cannot be deleted.
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.

Procedure

This section uses Microsoft SQL Server to MySQL synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-293 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 3-364 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-294 Synchronization instance details

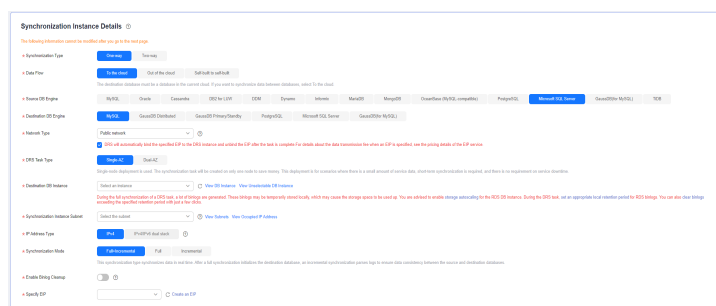


Table 3-365 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select MySQL .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	Select an RDS for MySQL instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	<p>Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-295 AZ



Table 3-366 Task AZ

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-296 AZ</p>

- Enterprise Project and Tags

Figure 3-297 Enterprise Project and Tags




Table 3-367 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.
- Source database information

Figure 3-298 Source database information

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS **RDS DB Instance**

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-368 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

- Destination database information

Figure 3-299 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-369 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MySQL instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary.
SSL Connection	If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-300 Synchronization Mode

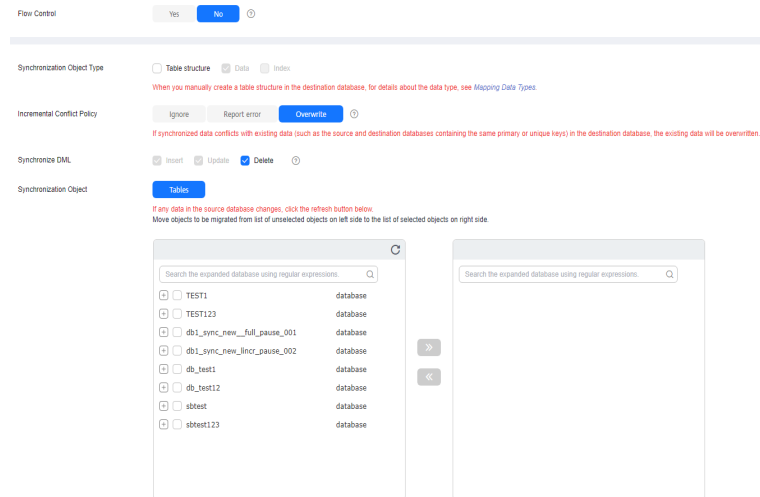
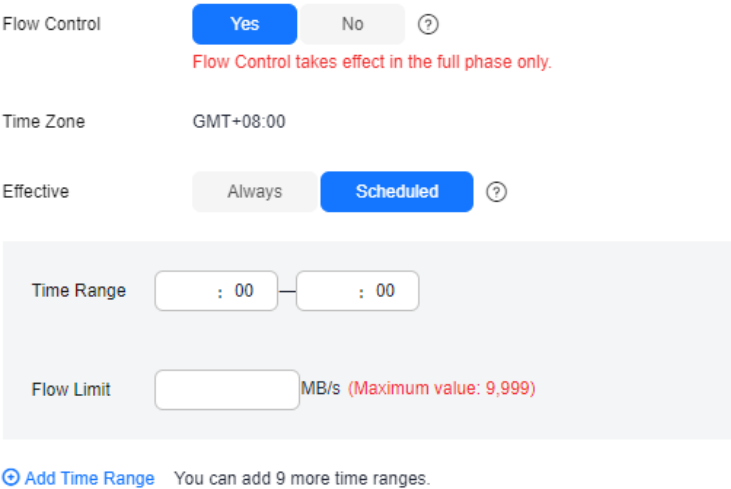



Table 3-370 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-301 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Parameter	Description
Start Point	<p>This option is available if you select Incremental in 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>If not specified, the current LSN of the database is used. For SQL Server, the value is the maximum log sequence number in the start_lsn column in the cdc.lsn_time_mapping system catalog.</p> <p>If specified, the valid LSN range is [minimum LSN reserved for CDC for the table to be synchronized, current LSN of the database]. Data changes after the LSN is specified are synchronized.</p> <p>1. Run the following command to query the minimum LSN reserved for CDC for the table to be synchronized:</p> <pre>select min(start_lsn) as minLsn from cdc.change_tables where source_object_id in (OBJECT_ID('[schemaXXX].[tableXXX]'));</pre> <p>If the data capture instance is not found or the caller is not authorized to access the changed data associated with the data capture instance, 0x00000000000000000000 is returned.</p> <p>2. Run the following command to query the current LSN of the database:</p> <pre>select sys.fn_cdc_get_max_lsn ();</pre> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-302 Task startup settings





Table 3-371 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.31 From Microsoft SQL Server to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-372 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	GaussDB(for MySQL) Primary/Standby

 NOTE

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 3-373](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-373 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: HIERARCHYID, TABLE, SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - The following functions can be used as default values during table structure synchronization: now, newid, getutcdate, and getdate. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> 1. The default value may be left empty. 2. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. - A database or table name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized. - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none">- DML statements, including INSERT, UPDATE, and DELETE, are supported.- DDL statements are not supported.- In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-374](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

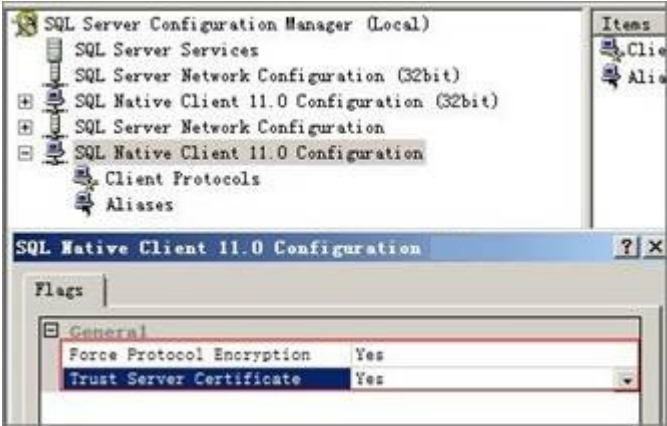
Table 3-374 Database account permission

Type	Full+Incremental Synchronization
Source database user	At least the sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	The user must have the SELECT, CREATE, DROP, INSERT, DELETE, UPDATE, ALTER, REFERENCES and INDEX permissions.

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-375 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - SQL Server Agent must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-303. <p style="text-align: center;">Figure 3-303 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: [] ? , . - The SQL Server character type uses the number of bytes as the length, and the GaussDB(for MySQL) character type uses the number of characters as the length and the character set is utf8mb4. Therefore, the number of bytes occupied by a field in the destination database is n x 4. - The length of each index cannot exceed 3,072 bytes after being converted into a field in the destination database. - The total number of indexes cannot exceed 64. - The total length of fields after being mapping to the destination database cannot exceed 66,535 bytes. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database must be a primary/standby GaussDB(for MySQL) instance. - The destination DB instance must have sufficient storage space.

Type	Constraints
	<ul style="list-style-type: none"> - The destination DB instance cannot contain databases with the same name as the source database (except system databases). • Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. - Do not use an imprecise value, such as floating-point type and date type, as a primary key or an equivalent condition after WHERE in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data type mapping, see Mapping Data Types from Microsoft SQL Server to GaussDB(for MySQL). - Object names will be converted to lowercase letters after being synchronized to the destination database. Therefore, the selected source database tables cannot contain tables with the same name but different letter cases. Otherwise, the synchronization fails. - Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. DDL operations performed on the source database will not be synchronized to the destination database. The IMAGE, TEXT, and NTEXT big data types cannot be deleted.
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-304 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description

0/256

Table 3-376 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-305 Synchronization instance details

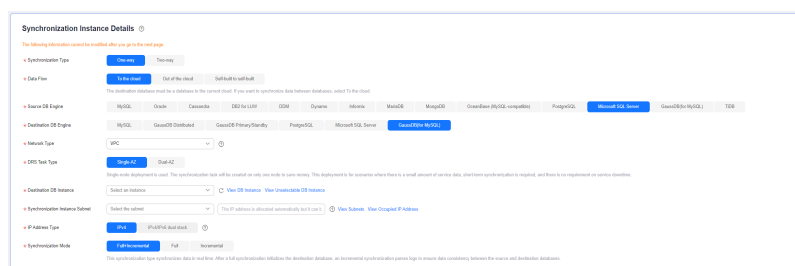


Table 3-377 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB(for MySQL) .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	Select a GaussDB(for MySQL) instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-306 AZ





Table 3-378 Task AZ

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-307 AZ</p> 

- Enterprise Project and Tags

Figure 3-308 Enterprise Project and Tags




Table 3-379 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.
- Source database information

Figure 3-309 Source database information

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS **RDS DB Instance**

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-380 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

- Destination database information

Figure 3-310 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-381 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-311 Synchronization Mode

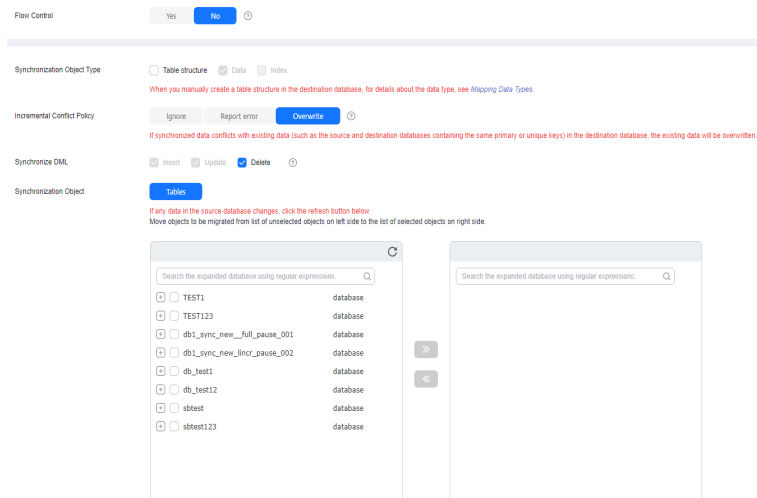
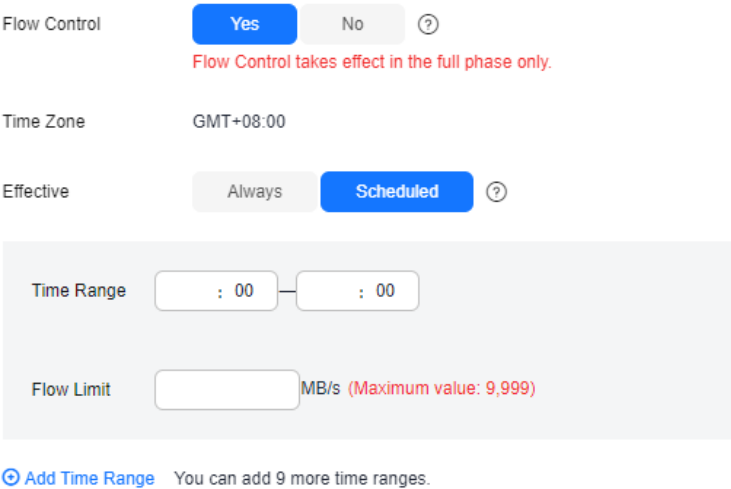



Table 3-382 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-312 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none">• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).• The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.)• For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

Parameter	Description
Start Point	<p>This option is available if you select Incremental in 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>If not specified, the current LSN of the database is used. For SQL Server, the value is the maximum log sequence number in the start_lsn column in the cdc.lsn_time_mapping system catalog.</p> <p>If specified, the valid LSN range is [minimum LSN reserved for CDC for the table to be synchronized, current LSN of the database]. Data changes after the LSN is specified are synchronized.</p> <p>1. Run the following command to query the minimum LSN reserved for CDC for the table to be synchronized: <pre>select min(start_lsn) as minLsn from cdc.change_tables where source_object_id in (OBJECT_ID('[schemaXXX].[tableXXX]'));</pre> If the data capture instance is not found or the caller is not authorized to access the changed data associated with the data capture instance, 0x00000000000000000000 is returned.</p> <p>2. Run the following command to query the current LSN of the database: <pre>select sys.fn_cdc_get_max_lsn ();</pre> For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-313 Task startup settings





Table 3-383 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.32 From Microsoft SQL Server to PostgreSQL

Supported Source and Destination Databases

Table 3-384 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) Microsoft SQL Server-compatible databases on other clouds (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 and Standard Edition 2016 SP2 or later, 2017, 2019 and 2022) 	RDS for PostgreSQL (versions 9.5, 9.6, 10, 11, 12, 13, 14 and 15)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-385 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-385 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● Supported field types: TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, SMALLDATETIME, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: HIERARCHYID, TABLE, SQL_VARIANT, GEOMETRY, and GEOGRAPHY ● Scope of full synchronization <ul style="list-style-type: none"> - Table structures, data, and indexes of selected tables can be synchronized. - A database or table name can contain a maximum of 58 characters, including only letters, digits, underscores (_), and hyphens (-). - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - During table structure synchronization in a full synchronization, an auto-increment column in the source database is synchronized to the destination database as a common column, and the auto-increment attribute is not synchronized. - For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. - If the default values contain values or functions that are not supported by the destination database, the table structure fails to be migrated. - In a full synchronization, calculated columns in the source database can be synchronized to the destination database as common columns. ● Scope of incremental synchronization <ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported.

Type	Precautions
	<ul style="list-style-type: none"> - In an incremental synchronization, calculated columns in the source database cannot be synchronized.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-386](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

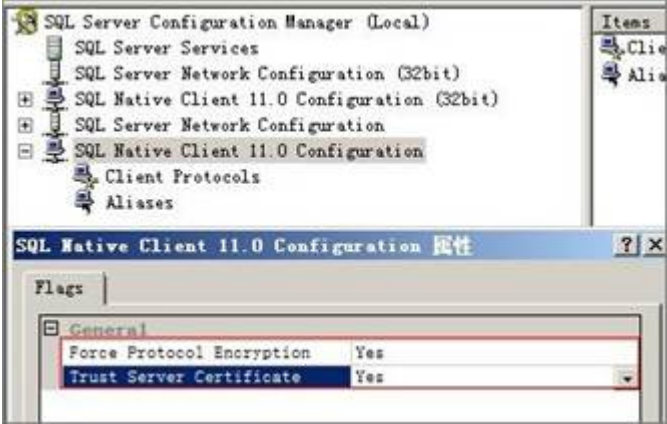
Table 3-386 Database account permission

Type	Full Synchronization	Incremental and Full+Incremental
Source database user	Sysadmin permission; db_datareader or db_owner permission for a database to be synchronized	At least the sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized
Destination database user	<p>The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.</p> <p>The account of the RDS for PostgreSQL instance has the preceding permissions by default.</p>	

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-387 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - SQL Server Agent must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 3-314. <p style="text-align: center;">Figure 3-314 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 58 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []?., - The SQL Server character type uses the number of bytes as the length, and the PostgreSQL character type uses the number of characters as the length and the character set is utf8mb4. Therefore, the number of bytes occupied by a field in the destination database is n x 4. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination DB instance must have sufficient storage space. - The destination DB instance cannot contain databases with the same name as the source database (except system databases). ● Other notes:

Type	Constraints
	<ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. - Do not use an imprecise value, such as floating-point type and date type, as a primary key or an equivalent condition after WHERE in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data type mapping, see Mapping Data Types from Microsoft SQL Server to PostgreSQL. - Some SQL Server databases (such as Azure SQL) do not support the use syntax and do not support synchronization of multiple databases. If you need to synchronize multiple databases, create multiple synchronization tasks.
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. DDL operations are not supported. Do not perform DDL operations on the source database. The IMAGE, TEXT, and NTEXT big data types cannot be deleted.
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-315 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description

0/256

Table 3-388 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-316 Synchronization instance information

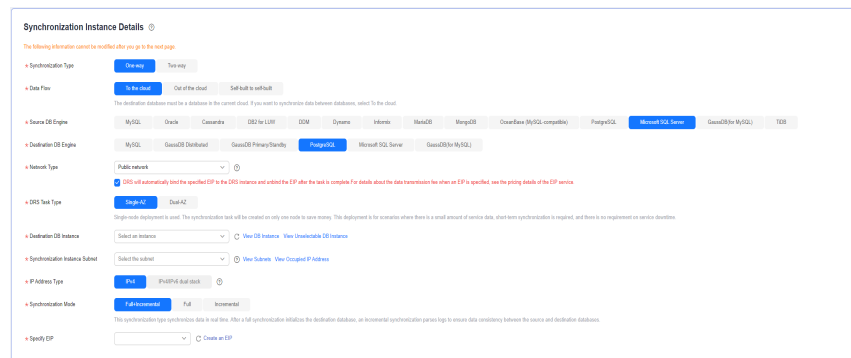


Table 3-389 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select PostgreSQL .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	<p>Select an RDS for PostgreSQL instance you created.</p>

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-317 AZ





Table 3-390 Task AZ

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-318 AZ</p> 

- Enterprise Project and Tags

Figure 3-319 Enterprise Project and Tags




Table 3-391 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.
- Source database information

Figure 3-320 Source database information

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

Database Type: Self-built on ECS RDS DB Instance

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Database Username:

Database Password:

Table 3-392 Source database settings

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance.
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

- Destination database information

Figure 3-321 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-393 Destination database settings

Parameter	Description
DB Instance Name	The RDS for PostgreSQL instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-322 Synchronization Mode

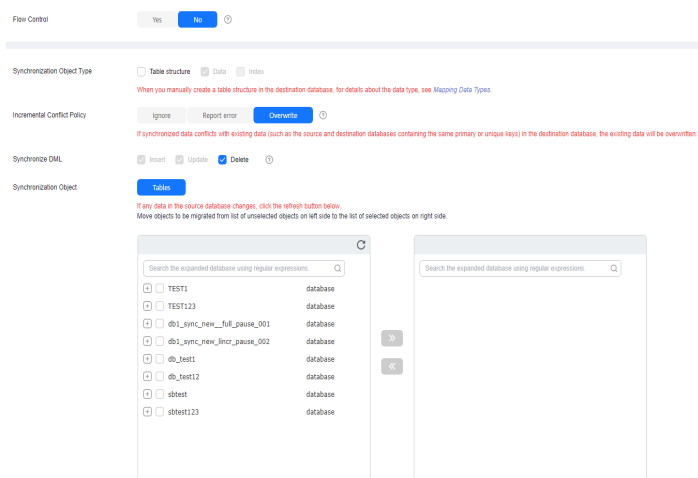
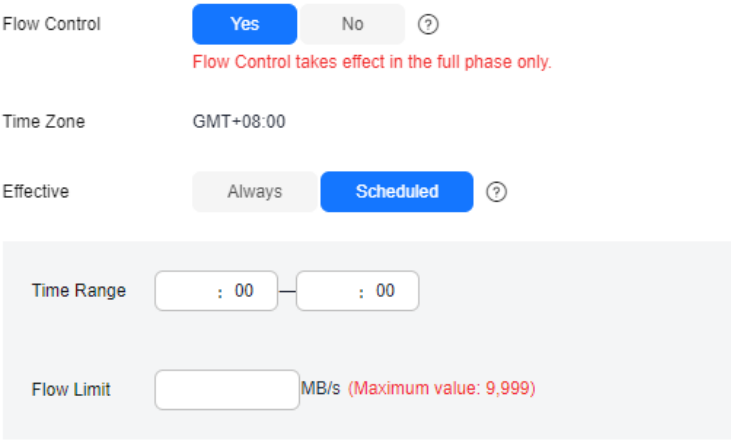



Table 3-394 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-323 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none">• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).• The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.)• For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

Parameter	Description
Start Point	<p>This option is available if you select Incremental in 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>If not specified, the current LSN of the database is used. For SQL Server, the value is the maximum log sequence number in the start_lsn column in the cdc.lsn_time_mapping system catalog.</p> <p>If specified, the valid LSN range is [minimum LSN reserved for CDC for the table to be synchronized, current LSN of the database]. Data changes after the LSN is specified are synchronized.</p> <p>1. Run the following command to query the minimum LSN reserved for CDC for the table to be synchronized:</p> <pre>select min(start_lsn) as minLsn from cdc.change_tables where source_object_id in (OBJECT_ID('[schemaXXX].[tableXXX]'));</pre> <p>If the data capture instance is not found or the caller is not authorized to access the changed data associated with the data capture instance, 0x00000000000000000000 is returned.</p> <p>2. Run the following command to query the current LSN of the database:</p> <pre>select sys.fn_cdc_get_max_lsn ();</pre> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-324 Task startup settings





Table 3-395 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.33 From MongoDB to DDS

Supported Source and Destination Databases

Table 3-396 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, 4.4 and 5.0) <p>NOTE The source database cannot be a GeminiDB Mongo instance. DDS 5.0 supports replica sets only. If the source database is a DDS DB instance, the source DB engine is DDS. Otherwise, the source DB engine is MongoDB-DDS.</p>	<p>DDS DB instances (versions 3.4, 4.0, 4.2, 4.4, and 5.0)</p> <p>NOTE The destination database version must be the same as or later than the source database version. DDS 5.0 supports replica sets only.</p>

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-397](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-397 Database account permissions

Type	Full+Incremental Synchronization
Source database user	<p>Replica set: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database.</p> <p>Single node: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database.</p>
Destination database user	The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-398](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-398 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. • System databases (such as local, admin, and config) and system collections cannot be synchronized. If service data is in a system database, run renameCollection to move the service data to the user database. • The statement for creating a view cannot contain a regular expression. • Collections that contain the _id field without indexes are not supported. • The first parameter of BinData() cannot be 2. • If the source database is a replica set instance of version 4.4 or later and the destination database version is earlier than 5.0, composite hash indexes are not supported. If the destination database version is 5.0, composite hash indexes are supported. • Time series collections are not supported in both the full and incremental phases. • Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-399](#).

Table 3-399 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• During the incremental synchronization, the Olog of the source database must be enabled.• If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is three days.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During task startup or full migration, you are not advised to perform DDL operations on the source database, such as deleting databases, collections, indexes, documents, or views. Otherwise, the synchronization may fail. ● During the synchronization, data rollback caused by a primary/standby switchover of the source database is not supported. ● Documents larger than 16 MB in the source database cannot be inserted or updated during full or incremental synchronization. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● If the source is a replica set instance, the following operations and commands are supported during incremental synchronization: <ul style="list-style-type: none"> - Adding, deleting, and updating documents - Creating and deleting collections - Creating and deleting indexes - Creating and deleting views - convertToCapped, collMod, and renameCollection commands ● In the incremental synchronization phase, concurrent replay is performed at the collection level to maintain the synchronization performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. <p>In either of the preceding scenarios, the task delay may increase.</p> <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● Replica set: The MongoDB replica set instance must be available and have primary nodes. ● The value of block_compressor is determined by stats().wiredTiger.creationString.block_compressor of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be synchronized. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on net.compression.compressors of the destination database. If the destination database version is DDS 4.2, DRS does not synchronize compression parameters because the destination database does not support compression parameter settings. ● If the MongoDB service of the source database is deployed with other services on the same server, set the value of the cacheSizeGB parameter to the half of the minimum idle cache for the WiredTiger engine of the source database. ● During an incremental synchronization of collections, renaming collections is not recommended. ● Specify the information about all primary and secondary nodes of the replica set to prevent the primary/standby switchover from affecting the synchronization task. If you enter information about multiple primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● To accelerate the synchronization, delete unnecessary indexes from the source database and retain only necessary indexes before the synchronization. You are advised not to create indexes for the source database during the synchronization. If indexes must be created, create them in the background. ● To prevent loopback, do not start tasks that migrate the same database to and out of the cloud at the same time. <p>NOTE The source and destination databases in a synchronization task can be the same database. To avoid replication loop, you must rename object names.</p>

Procedure

This section uses DDS replica set as an example to describe how to configure a MongoDB to DDS replica set synchronization task over a VPC network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-325 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: ..
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: ..

* Task Name: ⓘ

Description: ⓘ
0/256

Table 3-400 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-326 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow: Out of the cloud Self-hosted to self-hosted

Source DB Engine: MySQL DRZ for LUW MariaDB PostgreSQL GaussDB for MySQL

Destination DB Engine:

Network Type: ⓘ
 DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

Destination DB Instance: ⓘ

Synchronization Instance Subnet: ⓘ

Synchronization Mode: ⓘ
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Source DB Instance Type:

Specify EP: ⓘ

Table 3-401 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select MongoDB .
Destination DB Engine	Select DDS .
Network Type	<p>Available options: VPC, VPN or Direct Connect, and Public network. By default, the value is Public network.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you have created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <p>Full+Incremental: This synchronization type allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases.</p> <p>NOTE</p> <p>If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>

Parameter	Description
Source DB Instance Type	If Synchronization Mode is set to Full+Incremental , set this parameter based on the source database. Currently, the source database supports only replica set instances and single node instances with Olog enabled, which is a non-cluster database.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 3-327 Task type




Table 3-402 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-328 Enterprise Project and Tags




Table 3-403 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 3-329 Source database information

Source Database

Database Type: self-built database **DDS DB Instance**

DB Instance Name: [View DB Instance](#) [View Unselectable DB Instance](#)

Authentication Database:

Database Username:

Database Password:

SSL Connection:

Table 3-404 Source database settings

Parameter	Description
Database Types	The source database type can be Self-built on ECS or a DDS DB Instance . In this example, DDS DB Instance is selected.
VPC	This option is available if the source database type is Self-built on ECS . VPC is used to isolate networks for different services.
Subnet	This option is available if the source database type is Self-built on ECS . Subnet provides dedicated and isolated networks for different services and controls access. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database. This option is available if the source database type is Self-built on ECS .
DB Instance Name	The DDS DB instance as the source. This option is available if the source database type is DDS DB Instance .
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration


Figure 3-330 Destination database information

Destination Database

DB Instance Name

Authentication Database

Database Username

Database Password 

SSL Connection

Table 3-405 Destination database settings

Parameter	Description
DB Instance Name	The default value is the DB instance selected when the synchronization task is created and cannot be changed.
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-331 Synchronization objects

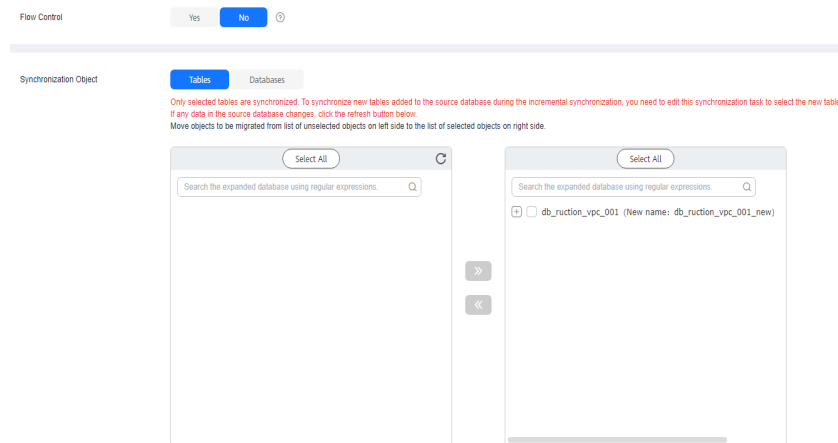
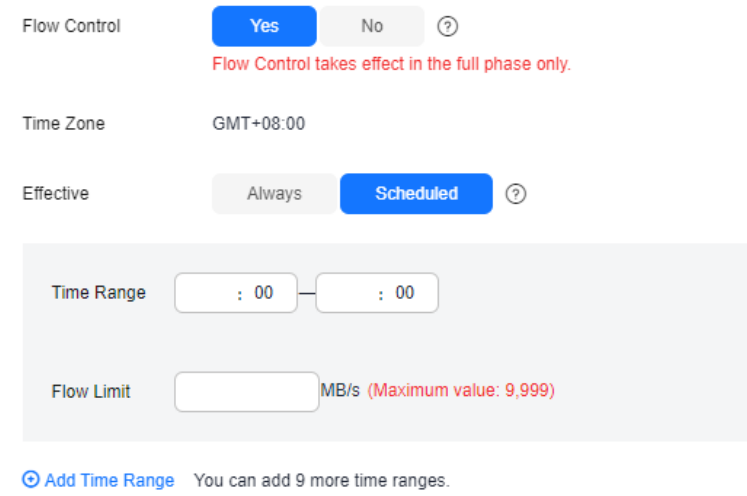




Table 3-406 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-332 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). The mapped name can contain 1 to 63 characters. The following characters are not allowed \. "\$ < ></p> <p>If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-333 Task startup settings





Table 3-407 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.34 From MariaDB to MariaDB

Supported Source and Destination Databases

Table 3-408 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MariaDB 10.3, 10.4, and 10.5 ECS-hosted MariaDB 10.3, 10.4, and 10.5 Other cloud MariaDB 10.3, 10.4, and 10.5 	RDS for MariaDB 10.5

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-409](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-409 Database account permission

Type	Full+Incremental
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)
Destination database user	The root account of RDS for MariaDB has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX

Supported Synchronization Objects

[Table 3-410](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-410 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-411](#).

Table 3-411 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: <code>.'<'>/\</code>• The source table and view names cannot contain non-ASCII characters, or the following characters: <code>.'<'>/\</code>• The column names in the source database tables cannot end with a backslash (<code>\</code>).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• The source and destination databases cannot contain tables that have the same names but do not have primary keys.• If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.• If the source database version is earlier than 10.4.3 and database-level synchronization is selected, after the task is started, do not use table creation statements (for example, create table t1 (c1 json)) containing JSON data type to create table structures or use column adding statements (for example, alter table t1 add column c1 json) containing JSON data type in the database selected for the synchronization object. Otherwise, data will be inconsistent.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. During data synchronization, do not upgrade the source MariaDB database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MariaDB database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. Data inconsistency may occur when the MyISAM table is modified during a full synchronization. <p>Incremental synchronization</p> <ul style="list-style-type: none"> DDL statements can be synchronized during incremental synchronization. Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. You can add additional objects in the incremental synchronization phase. <p>Synchronization comparison</p> <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none"> If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> – If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. – If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. – You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination databases are the same RDS DB instance, real-time synchronization without database mapping is not supported. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● The source database does not support point-in-time recovery (PITR). ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common

Type	Restrictions
	<p>index, the table structure may fail to be created. You are advised to use a unique index.</p> <ul style="list-style-type: none"> • Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. • The partitioned table does not support column mapping. • After a task is created, the destination database cannot be set to read-only. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

This section describes how to use DRS to configure a MariaDB synchronization task over a public network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-334 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-412 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-335 Synchronization instance information

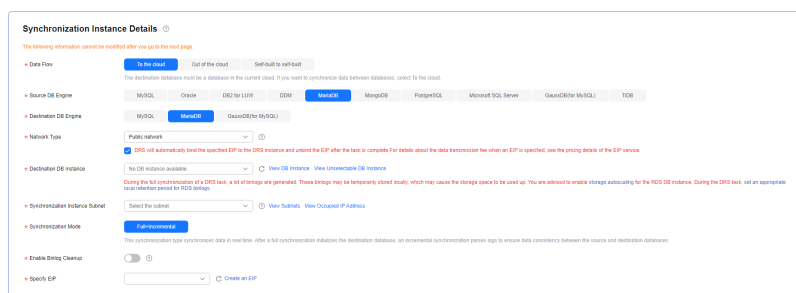


Table 3-413 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select MariaDB .
Destination DB Engine	Select MariaDB .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is the default value and is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	The RDS for MariaDB instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <p>Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases.</p> <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>

Parameter	Description
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 3-336 Task type




Table 3-414 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-337 Enterprise Project and Tags




Table 3-415 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 3-338 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 3-416 Source database information

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting , Full synchronization , Incremental synchronization , or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-339 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

✔ Test successful

Table 3-417 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MariaDB instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	Password for logging in to the destination database.
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-340 Synchronization objects

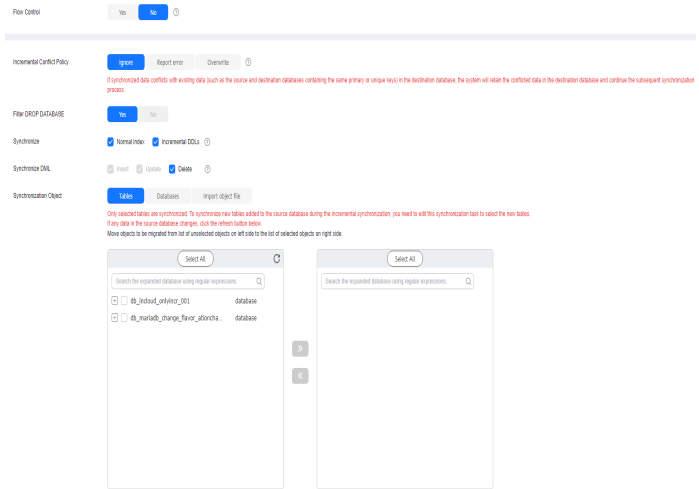
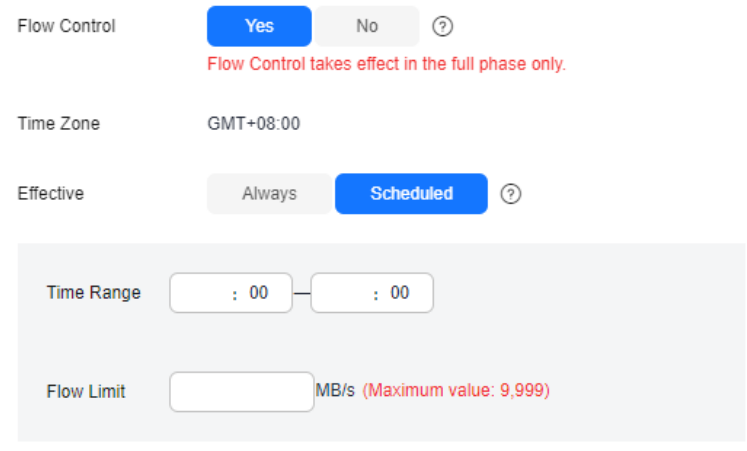



Table 3-418 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-341 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**. For details about how to configure related rules, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-342 Task startup settings





Table 3-419 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none">You can set this parameter only for pay-per-use tasks.Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.35 From MariaDB to MySQL

Supported Source and Destination Databases

Table 3-420 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 	RDS for MySQL 5.5, 5.6, 5.7, and 8.0 NOTE <ul style="list-style-type: none"> If the source database version is MariaDB 10.0 or 10.1, you are advised to select MySQL 5.6 or later as the destination database. If the source database version is MariaDB 10.2, 10.3 or 10.4, you are advised to select MySQL 5.7 or later as the destination database. If the source database version is MariaDB 10.5, you are advised to select MySQL 8.0 as the destination database.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-421](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-421 Database account permission

Type	Incremental and Full+Incremental
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)

Type	Incremental and Full+Incremental
Destination database user	The root account of RDS for MySQL has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Supported Synchronization Objects

Table 3-422 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-422 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see **Pre-check Items**. In addition to the pre-check items, you need to pay attention to the items listed in **Table 3-423**.

Table 3-423 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: .'<>/\• The table and view names in the source database cannot contain non-ASCII characters, or the following special characters: <>/\• The column names in the source database tables cannot end with a backslash (\).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• The source and destination databases cannot contain tables that have the same names but do not have primary keys.• If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● During data synchronization, do not upgrade the source MariaDB database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MariaDB database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● You can add additional objects during an incremental synchronization. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being

Type	Restrictions
	<p>falsely reported and reduce the impact on the source database and DRS tasks.</p> <ul style="list-style-type: none">• During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent.• Data cannot be compared during full synchronization.• Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> – If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. – If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. – You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions? ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the source database does not have a view but has a table structure with the same name as a view, do not run the drop view if exists SQL statement on the source database. Otherwise, the task fails. When drop view if exists is executed in MariaDB, if the corresponding view is not found but there is a table with the same name as the view, no exception is reported and binlog records are generated. When the command is synchronized to the destination database and drop view if exists is executed in MySQL in the same case, the error message "is not view" is displayed. As a result, the task fails. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided:

Type	Restrictions
	<ul style="list-style-type: none"> <li data-bbox="627 297 1406 495">– Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. <li data-bbox="627 510 1422 707">– Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. <li data-bbox="588 723 1410 786">● The source database does not support point-in-time recovery (PITR). <li data-bbox="588 801 1366 864">● The destination database cannot be restored to a point in time when a full synchronization was being performed. <li data-bbox="588 880 1398 1010">● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. <li data-bbox="588 1025 1262 1088">● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. <li data-bbox="588 1104 1353 1135">● The partitioned table does not support column mapping. <li data-bbox="588 1151 1422 1214">● After a task is created, the destination database cannot be set to read-only. <li data-bbox="588 1229 1406 1520">● During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. <li data-bbox="588 1536 1398 1733">● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

This section describes how to use DRS to configure a real-time synchronization task from MariaDB to RDS for MySQL over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-343 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page. The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 3-424 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-344 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud Self-host to self-host

The replication instance must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

• Source DB Engine MySQL Oracle DB2 for LUW DB2 MongoDB PostgreSQL Microsoft SQL Server GaussDB(for MySQL) TDSE

• Destination DB Engine MySQL MariaDB GaussDB(for MySQL)

• Network Type Public network ⓘ

• DRS Task Type SingleAZ CrossAZ

Single-node deployment is used. This synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service durability.

• Destination DB Instance No DB instance available View DR Instance View Unavailable DB Instance

During the full synchronization of a DRS task, all DB storage are generated. Please design into the temporary disk capacity, which may cause the storage space to be used up. You are advised to enable storage autoresizing for the RDS DB instance. During the DRS task, set an appropriate local retention period for RDS storage. You can also clear strings according to the specified retention period with just a few clicks.

• Synchronization Instance Subnet Select the subnet View Subnets View Occupied IP Address

• Synchronization Mode Full-Initialization Incremental

The synchronization task synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization process begins to ensure data consistency between the source and destination databases.

• Enable Bring Cleanup ⓘ

• Selects EP Create an EP

Table 3-425 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select MariaDB .
Destination DB Engine	Select MySQL .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is the default value and is used as an example. <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task .

Parameter	Description
Destination DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-345 Specifications




Table 3-426 Task type information

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-346 AZ</p>

- Enterprise Project and Tags

Figure 3-347 Enterprise Project and Tags




Table 3-427 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 3-348 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 3-428 Source database information

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.

Parameter	Description
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-349 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-429 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MySQL instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	Password for logging in to the destination database.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-350 Synchronization objects

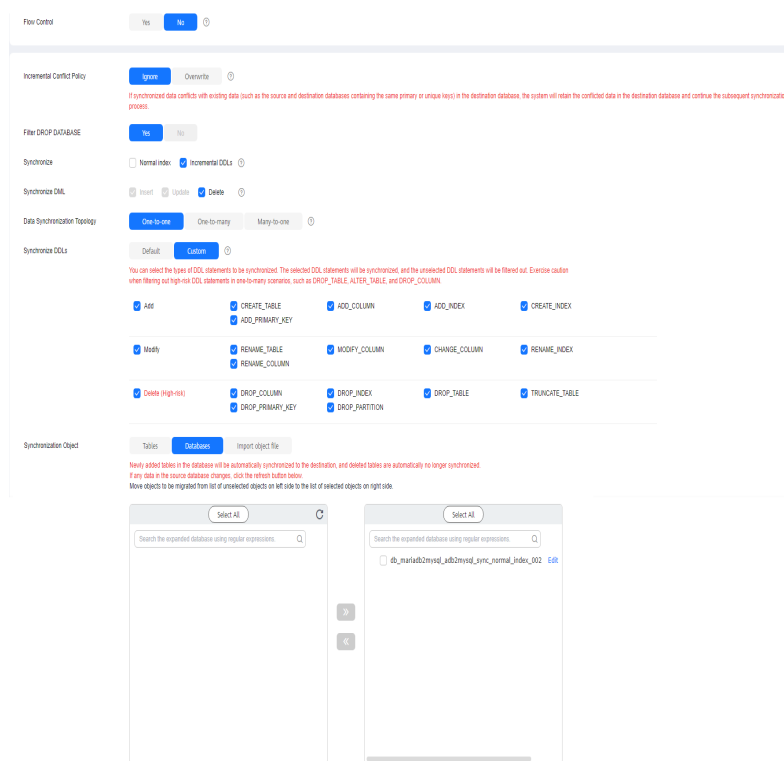
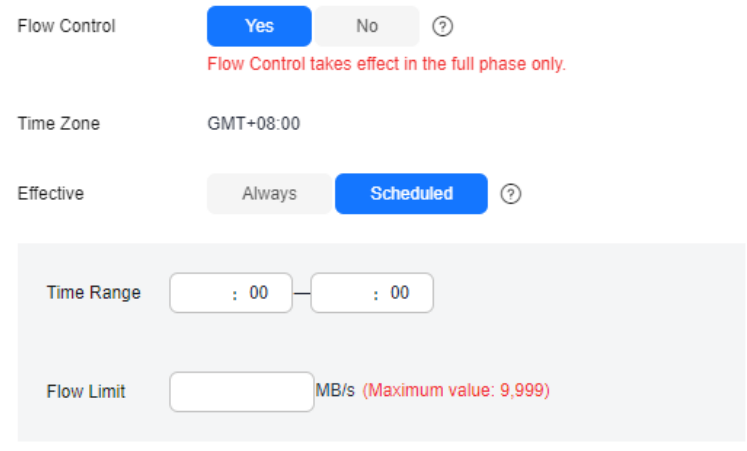



Table 3-430 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-351 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>

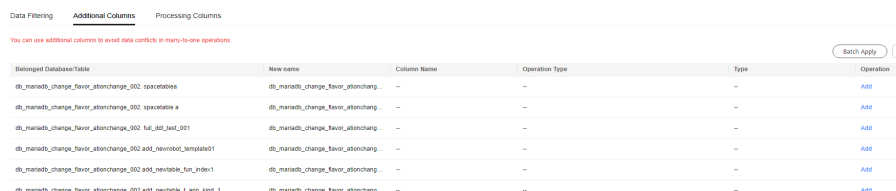
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-352 Processing data



Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
db_mansdb_change_favor_atonchange_002_tpacostable	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_tpacostable a	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_t66_db_text_001	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newrobot_template01	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newtable_for_index1	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newtable_1_app_kvml_1	db_mansdb_change_favor_atonchang	--	--	--	ADD

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-353 Task startup settings

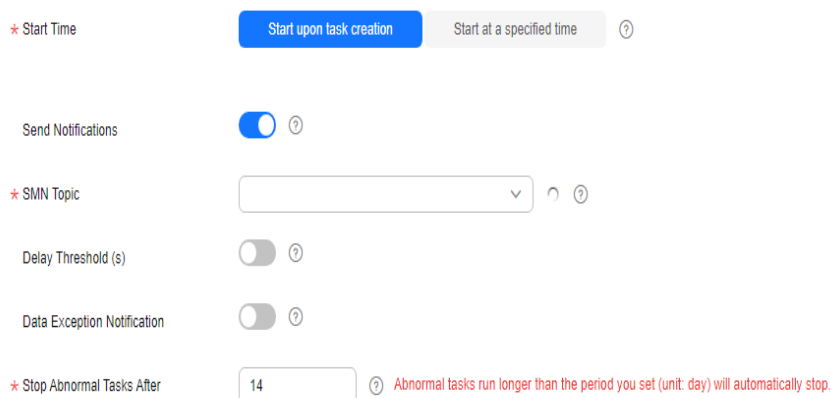



Table 3-431 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.36 From MariaDB to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-432 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5	GaussDB(for MySQL) Primary/Standby (version 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-433](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-433 Database account permission

Type	Incremental and Full+Incremental
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)
Destination database user	The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Supported Synchronization Objects

Table 3-434 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-434 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none">• Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Events and triggers cannot be synchronized.• Table names cannot be mapped for tables on which views, stored procedures, and functions depend.• When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see **Pre-check Items**. In addition to the pre-check items, you need to pay attention to the items listed in **Table 3-435**.

Table 3-435 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: .'<>/\• The table and view names in the source database cannot contain non-ASCII characters, or the following special characters: <>/\• The column names in the source database tables cannot end with a backslash (\).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• The source and destination databases cannot contain tables that have the same names but do not have primary keys.• If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● During data synchronization, do not upgrade the source MariaDB database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MariaDB database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● You can add additional objects during an incremental synchronization. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? <p>Synchronization comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being

Type	Restrictions
	<p>falsely reported and reduce the impact on the source database and DRS tasks.</p> <ul style="list-style-type: none">• During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent.• Data cannot be compared during full synchronization.• Do not limit the synchronization speed during data comparison. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> – If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. – If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. – You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the source database does not have a view but has a table structure with the same name as a view, do not run the drop view if exists SQL statement on the source database. Otherwise, the task fails. When drop view if exists is executed in MariaDB, if the corresponding view is not found but there is a table with the same name as the view, no exception is reported and binlog records are generated. When the command is synchronized to the destination database and drop view if exists is executed in GaussDB(for MySQL) in the same case, the error message "is not view" is displayed. As a result, the task fails. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the

Type	Restrictions
	<p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The source database does not support point-in-time recovery (PITR). • The destination database cannot be restored to a point in time when a full synchronization was being performed. • Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. • Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. • The partitioned table does not support column mapping. • After a task is created, the destination database cannot be set to read-only. • During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

This section describes how to use DRS to configure a real-time synchronization task from MariaDB to GaussDB(for MySQL) over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-354 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] .. v
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu] .. v

Task Name: DRS-5678 ⓘ

Description: [Text area] ⓘ
0/256

Table 3-436 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-355 Synchronization instance details

Synchronization Instance Details ⓘ

⚠ The following information cannot be modified after you go to the next page.

- Data Flow:** [To the cloud] [Out of the cloud] [Self-built to self-built]

The destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.
- Source DB Engine:** MySQL, Oracle, DRS for LUW, ODM, **MaxCompute**, MongoDB, PostgreSQL, Microsoft SQL Server, GaussDB for MySQL, TDS
- Destination DB Engine:** MySQL, MaxCompute, **GaussDB for MySQL**
- Network Type:** Public network [Dropdown menu] ⓘ

DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.
- DRS Task Type:** Single-AZ, **Dual-AZ**

Single node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime.
- Destination DB Instance:** [No DB instance available] [Dropdown menu] ⓘ [View DB Instance](#) [View Instances DB Instance](#)
- Synchronization Instance Subnet:** [Select the subnet] [Dropdown menu] ⓘ [View Subnets](#) [View Occupied IP Address](#)
- Synchronization Mode:** **Full-overwrite**, Incremental

This synchronization type synchronizes data in real time, after a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
- Specify EIP:** [Dropdown menu] ⓘ [Create an EIP](#)

Table 3-437 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select MariaDB .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is the default value and is used as an example. <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task .

Parameter	Description
Destination DB Instance	The GaussDB(for MySQL) instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-356 Specifications

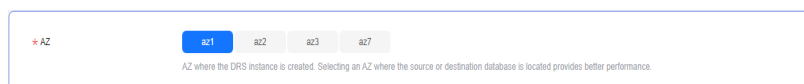



Table 3-438 Task type information

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-357 AZ</p> 

- Enterprise Project and Tags

Figure 3-358 Enterprise Project and Tags




Table 3-439 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 3-359 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 3-440 Source database information

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.

Parameter	Description
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-360 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-441 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	Password for logging in to the destination database.
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-361 Synchronization objects

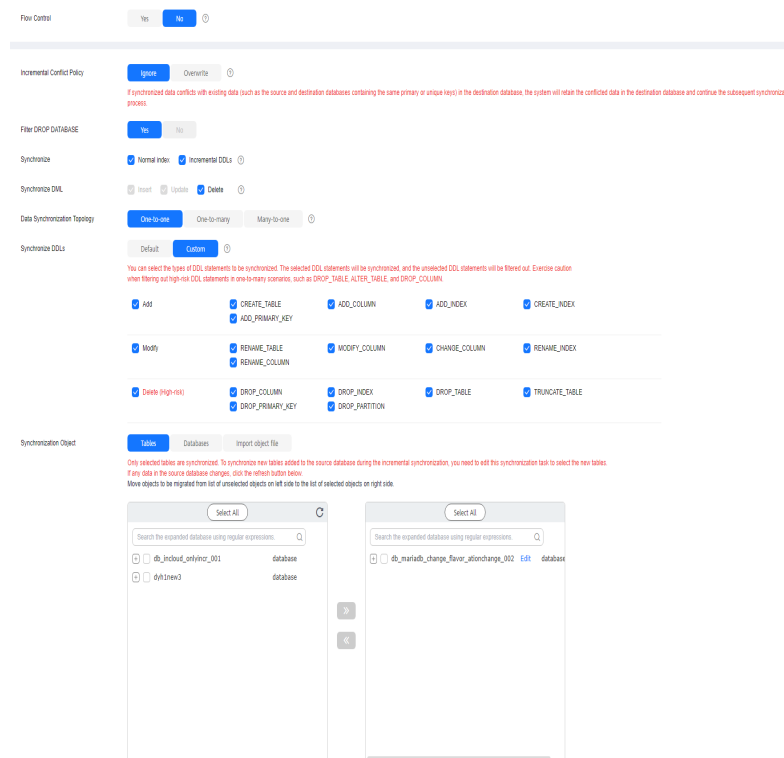
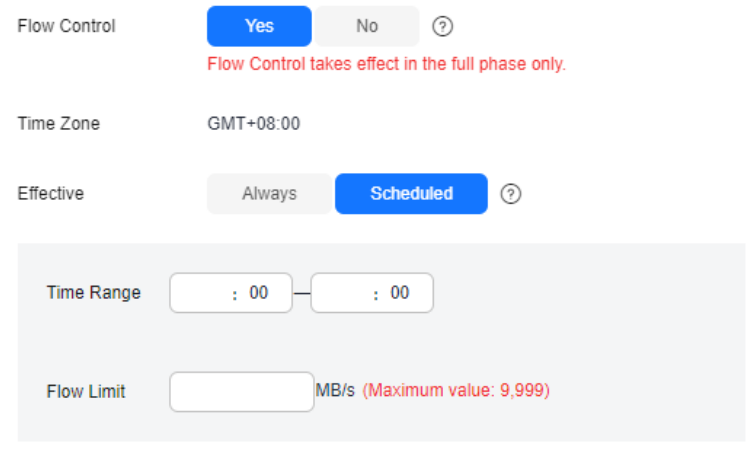



Table 3-442 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-362 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>

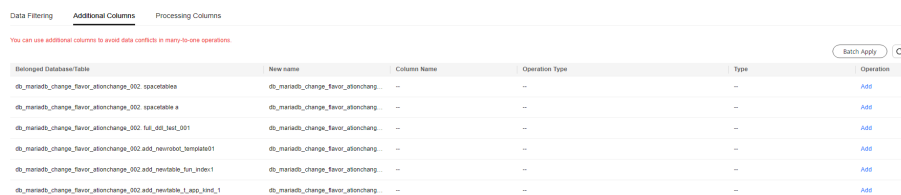
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-363 Processing data



Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
db_mansdb_change_favor_atonchange_002_tpacostable	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_tpacostable a	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_t66_db_text_001	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newrobot_template01	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newtable_for_index1	db_mansdb_change_favor_atonchang	--	--	--	ADD
db_mansdb_change_favor_atonchange_002_at6_newtable_1_app_kvml_1	db_mansdb_change_favor_atonchang	--	--	--	ADD

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-364 Task startup settings

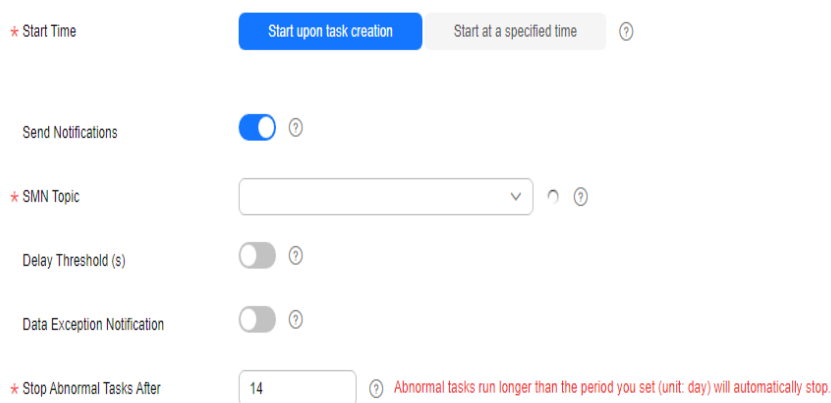



Table 3-443 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.37 From GaussDB(for MySQL) to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-444 Supported databases

Source DB	Destination DB
GaussDB(for MySQL) Primary/Standby (version 8.0)	GaussDB(for MySQL) Primary/Standby (version 8.0)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-445](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-445 Database account permissions

Type	Incremental and Full+Incremental
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT
Destination database user	The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-446 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-446 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Events and triggers cannot be synchronized.• Table names cannot be mapped for tables on which views, stored procedures, and functions depend.• When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see **Pre-check Items**. In addition to the pre-check items, you need to pay attention to the items listed in **Table 3-447**.

Table 3-447 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: '<>\'\"• The source table and view names cannot contain non-ASCII characters, or the following characters: '<>\'\"• The column names in the source database tables cannot end with a backslash (\).• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• The source and destination databases cannot contain tables that have the same names but do not have primary keys.

Type	Restrictions
Restrictions on usage	<p>General</p> <p>During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.</p> <p>Full synchronization</p> <ul style="list-style-type: none"> • Data inconsistency may occur when the MyISAM table is modified during a full synchronization. • DDL operations are not supported during full synchronization. <p>Incremental synchronization</p> <p>During incremental synchronization, some DDL operations are supported.</p> <ul style="list-style-type: none"> • In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. • Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> – If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. – If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. – You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since GaussDB(for MySQL) 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*:180023 INVISIBLE */, PRIMARY KEY (`id`));</pre> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source and destination databases are the same GaussDB(for MySQL) instance, real-time synchronization without database mapping is not supported. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● The source database does not support point-in-time recovery (PITR).

Type	Restrictions
	<ul style="list-style-type: none"> • The destination database cannot be restored to a point in time when a full synchronization was being performed. • Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. • Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. • If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. • After a task is created, the destination database cannot be set to read-only. • If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. • During database name mapping, function objects, views, and stored procedures may fail to be created because they are referenced by a database table name before the mapping. DRS ignores the error reported during the creation of these objects in the full synchronization phase. As a result, these objects are lost in the destination database. In the incremental synchronization phase, if database name mapping is required, DDL operations (including CREATE, MODIFY, and DELETE) related to function objects, views, and stored procedures are not synchronized to the destination database.

Procedure

This section uses real-time synchronization from GaussDB(for MySQL) to GaussDB(for MySQL) as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-365 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Project

* Task Name

Description

0/256

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Table 3-448 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-366 Synchronization instance information

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: **To the cloud** (Selected) | Out of the cloud | Self-built to self-built

Source DB Engine: **GaussDB(for MySQL)** (Selected) | Oracle | DB2 for LUW | COS | MariaDB | MongoDB | PostgreSQL | Microsoft SQL Server | TDS

Destination DB Engine: **GaussDB(for MySQL)** (Selected)

Network Type: **Public network** (Selected) | Private network

DRS Task Type: **Single AZ** (Selected) | Dual-AZ

Destination DB Instance: **No DB instance available** (Selected) | View DB Instance | View Unavailable DB Instance

Synchronization Instance Subnet: **Select the subnet** (Selected) | View Subnets | View Occupied IP Address

Synchronization Mode: **Full-to-incremental** (Selected) | Incremental

Specify EIP: **None** (Selected) | Create an EIP

Table 3-449 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select GaussDB(for MySQL) .

Parameter	Description
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	An available GaussDB(for MySQL) instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 3-367 Task type





Table 3-450 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-368 AZ</p> 

- Enterprise Project and Tags

Figure 3-369 Enterprise Project and Tags




Table 3-451 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-370 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 3-452 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-371 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 3-453 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted.
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-372 Synchronization Mode

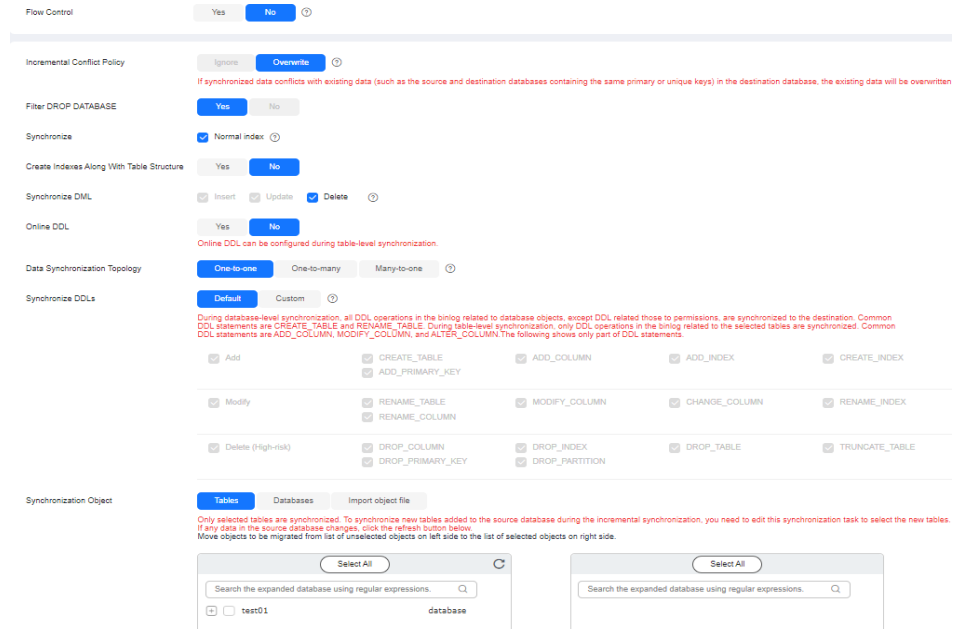
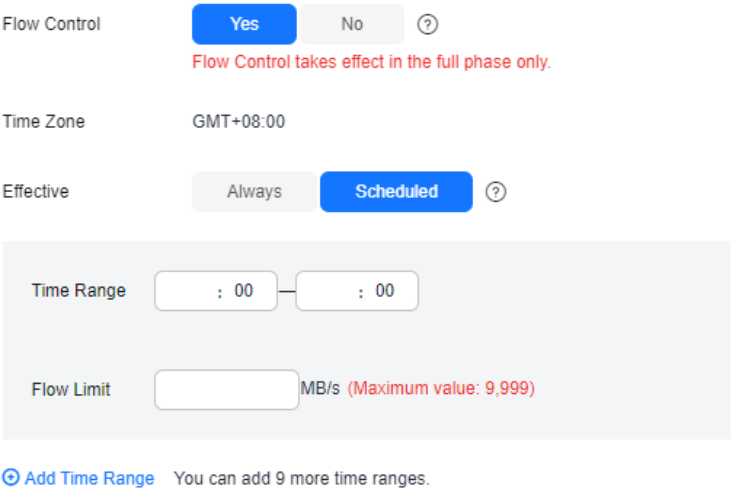


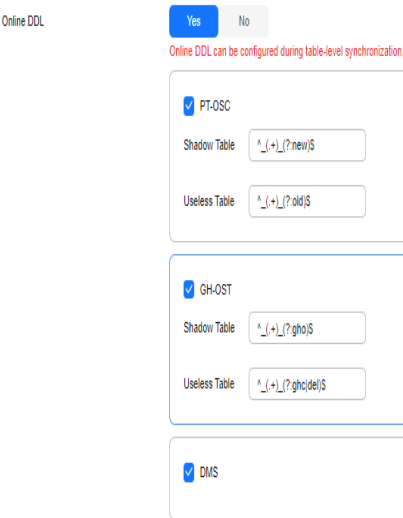
Table 3-454 Synchronization mode and object


Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-373 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Read on Standby Node	<p>If the source in Step 3 is an optional instance, data can be read from the standby database. You can enable or disable the function.</p> <ul style="list-style-type: none"> ● If enabled, full data is read from the standby database, and incremental binlogs are obtained from the primary database. <ul style="list-style-type: none"> – After this function is enabled, the primary and standby databases at the source end must be available, and the synchronization between the primary and standby databases is normal. If there are tables without primary keys in the synchronization object, the source database user must have the RELOAD permission. – After DDLs are executed on the primary database at the source end, the standby database disconnects the synchronization between the primary and standby databases. As a result, full synchronization of the DRS task may fail. If the table structure is not modified, you can resume the task. – If the primary/standby role at the source end changes, full synchronization of the DRS task may fail. In this case, you can resume the synchronization to restore the task. – Data can be read from the standby database only when the source is an optional instance. If the source is a self-managed database, the IP address cannot be entered. ● If disabled, both full data and incremental binlogs are obtained from the primary database.

Parameter	Description
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Whether to synchronize normal indexes.</p> <p>By default, DRS synchronizes the primary key or unique index. A normal index refers to an index other than the primary key or unique index. If you select Normal index, all indexes will be synchronized. If you do not select normal index, only the primary key or unique index will be synchronized.</p>
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none"> • Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase. • No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>

Parameter	Description
Synchronize DDLs	<p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OCS, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 3-374 Online DDL</p>  <p>The screenshot shows the 'Online DDL' configuration interface. At the top, there is a toggle switch for 'Online DDL' with 'Yes' selected. Below the toggle, a red message states 'Online DDL can be configured during table-level synchronization.' There are three sections for tool configuration:</p> <ul style="list-style-type: none"> PT-OSC: Checked. 'Shadow Table' field contains '^_(+)?_?new/\$'. 'Useless Table' field contains '^_(+)?_?old/\$'. GH-OST: Checked. 'Shadow Table' field contains '^_(+)?_?gho/\$'. 'Useless Table' field contains '^_(+)?_?gho(del)/\$'. DMS: Checked. No input fields are visible. No: Table-level synchronization does not support Online DDL synchronization.

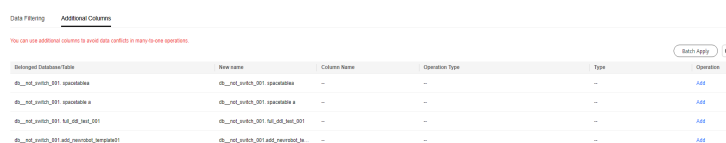
Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Processing Data** page, filter data or add additional columns for the table object to be processed, and click **Next**.

- If you need to set data filtering, click **Data filtering** and set related filtering rules.
- If you need to add additional columns, click the **Additional Columns** tab, click **Add** in the **Operation** column, and enter the column name and the operation type.

For details about related operations, see [Processing Data](#).

Figure 3-375 Processing data



Belonged Database Table	New name	Column Name	Operation Type	Type	Operation
db_001_mysql_001_saptable1	db_001_mysql_001_saptable1a	--	--	--	Add
db_001_mysql_001_saptable1a	db_001_mysql_001_saptable1a	--	--	--	Add
db_001_mysql_001_sap_001	db_001_mysql_001_sap_001	--	--	--	Add
db_001_mysql_001_sap_001	db_001_mysql_001_sap_001	--	--	--	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-376 Task startup settings

Table 3-455 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.38 From Cassandra to GeminiDB Cassandra

Supported Source and Destination Databases

Table 3-456 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Cassandra databases ECS-hosted Cassandra databases 	<ul style="list-style-type: none"> GeminiDB Cassandra instances

Supported Synchronization Objects

[Table 3-457](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-457 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> Supported scenario: Full synchronization Supported fields: ASCII, BIGINT, BLOB, BOOLEAN, COUNTER, DATE, DECIMAL, DOUBLE, DURATION, FLOAT, INET, INT, LIST, MAP, SET, SMALLINT, TEXT, TIME, TIMESTAMP, TIMEUUID, TINYINT, UUID, VARCHAR and VARINT. Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> During full synchronization, only table data and table structures can be synchronized. Table-level and column-level TTLs cannot be synchronized. Tables containing fields of the FROZEN, TUPLE, and USER-DEFINED types cannot be synchronized. If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. If you select synchronization objects by importing an object file, the token of a table cannot be specified. The default value 12 is used.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-458](#).

DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-458 Database user permissions

Type	Full Synchronization
Source database user	The user must have the following minimum permissions: SELECT permission on system catalogs system_auth.roles and system_auth.role_permissions and on the tables to be synchronized.
Destination database user	The user must have the following permissions: SELECT permission on system catalogs system_auth.roles and system_auth.role_permissions and CREATE permission on keyspaces. If the table to be synchronized already exists in the destination database, the SELECT and MODIFY permissions on the table are required.

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS full synchronization consists of three phases: task start, full synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-459 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> • The primary data center can be configured. You can enter a data center name. This configuration affects the load balancing policy and consistency level for data read from the source database. You can run the following command to query a data center name: select data_center from system.local; or select data_center from system.peers; • The default load balancing policy for data read is RoundRobinPolicy. When a primary data center is configured, the policy is DCAwareRoundRobinPolicy. If this configuration is incorrect, data may fail to be read. • The default consistency level for data read is EACH_QUORUM. In earlier versions (2.0 and 2.1), the level is ALL. When a primary data center is configured, the level is LOCAL_QUORUM. • The number of table tokens can be set during synchronization. The Z01_TABLE_TOKENS attribute can be specified when a table is created.

Type	Restrictions
Full synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not execute any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Procedure

This section uses real-time synchronization from Cassandra to GeminiDB Cassandra as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-377 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region .. v

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0256

Table 3-460 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-378 Synchronization instance information

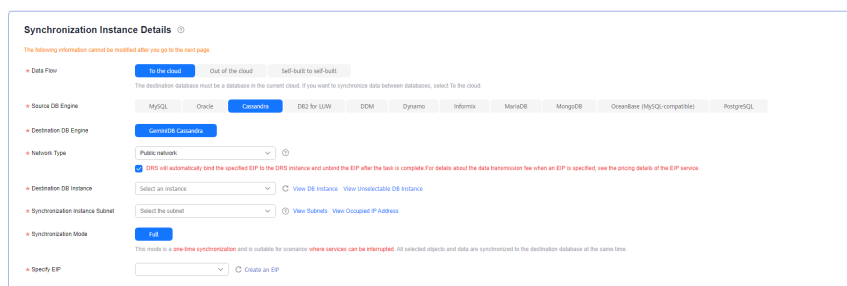


Table 3-461 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Cassandra .
Destination DB Engine	Select GeminiDB Cassandra .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
Destination DB Instance	Select an available GeminiDB Cassandra instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	The synchronization mode supported by a DRS task. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . Full: All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 3-379 AZ




Table 3-462 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-380 Enterprise Project and Tags




Table 3-463 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-381 Source database information

Source Database
System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 3-464 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 3-382 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password ⓘ

✔ Test successful

Table 3-465 Destination database settings

Parameter	Description
DB Instance Name	The GeminiDB Cassandra instance selected when you created the synchronization task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-383 Synchronization mode

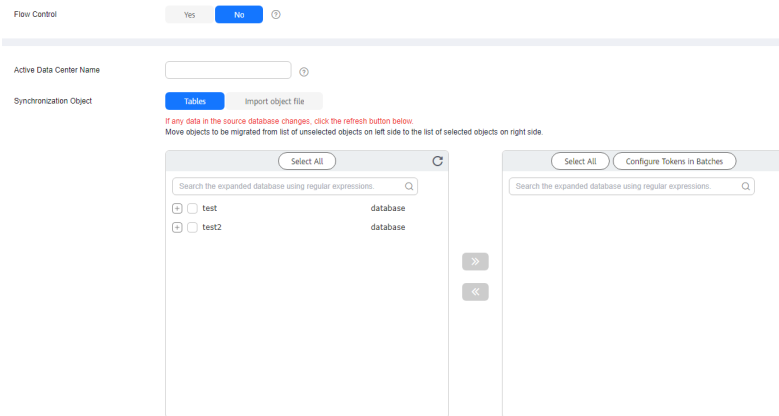
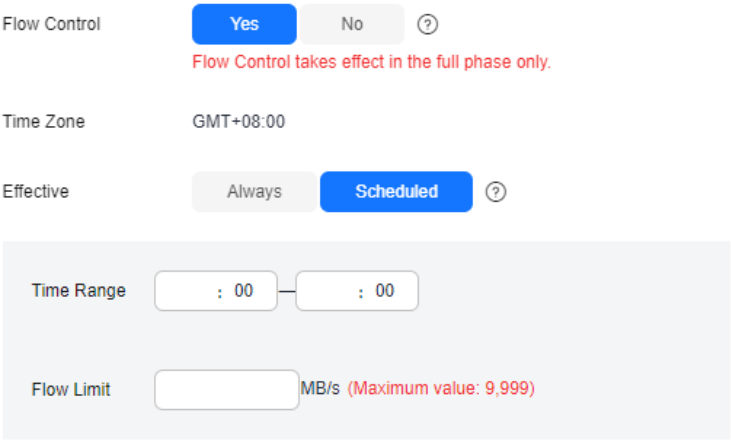



Table 3-466 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-384 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Active Data Center Name	<p>Enter a data center name. This configuration affects the load balancing policy and consistency level for data read from the source database. You can run the following command to query a data center name:</p> <pre>select data_center from system.local; or select data_center from system.peers;</pre>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file based on your service requirements.</p> <ul style="list-style-type: none"> On the right pane, you can set tokens for selected objects in batches. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). A mapped table name can contain 1 to 48 characters. Only letters, digits, and underscores (_) are allowed. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-385 Task startup settings

Table 3-467 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.39 From Dynamo to GeminiDB Cassandra

Supported Source and Destination Databases

Table 3-468 Supported databases

Source DB	Destination DB
DynamoDB on other clouds (web services)	GeminiDB Cassandra 3.11

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-469 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-469 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Supported scenarios: full synchronization and full +incremental synchronization. • Supported fields: String, Number, Binary, Boolean, Null, Map, List, String Set, Number Set, and Binary Set. • Table-level synchronization is supported. <ul style="list-style-type: none"> - During full synchronization, only table data and table structures can be synchronized. - TTL synchronization is not supported. - Counter tables cannot be synchronized. - All basic types of data in Map of the source table are converted into character strings in Map of the destination table, such as digits and binary data. The binary data is converted into hexadecimal character strings. If the source table has nested List, Map, and Set, the values in Map of the destination table are converted into JSON data. - All basic types of data in List of the source table are converted into character strings in List of the destination table, such as digits and binary data. The binary data is converted into hexadecimal character strings. If the source table has nested List, Map, and Set, the values in List of the destination table are converted into JSON data. - The token of a table can be specified during object selection.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-470](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-470 Database user permissions

Type	Full and Full+Incremental
Source database user	AmazonDynamoDBFullAccess grants access to all tables and features. You can also create a custom policy to grant only access to required resources. The user must have the following minimum permissions: DescribeTable, Scan, GetItem, DescribeStream, GetRecords, ListStreams, GetShardIterator, ListTables and DescribeEndpoints.
Destination database user	The user must have the following permissions: SELECT permission on system_auth.roles and system_auth.role_permissions, CREATE permission on keyspaces, and CREATE, SELECT, MODIFY, and ALTER permissions on synchronization tables.

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-471 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> The default read consistency level of the destination database is set to LOCAL_QUORUM. Table tokens can be set during synchronization. The Z01_TABLE_TOKENS attribute can be specified when a table is created.
Full synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> Do not change the port of the destination database, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Procedure

This section uses real-time synchronization from DynamoDB to GeminiDB Cassandra as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-386 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ
0/256

Table 3-472 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-387 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Synchronization Type One-way Two-way

• Data Flow To the cloud Out of the cloud Self-kull to self-kull
This destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

• Source DB Engine

• Destination DB Engine

• Network Type ⓘ
 DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance ⓘ

• Synchronization Instance Subnet ⓘ

• IP Address Type IPv4 IPv6 IPv4/IPv6 dual stack ⓘ

• Synchronization Mode Full-Initialization Full
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Specify EIP ⓘ

Table 3-473 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
Source DB Engine	Select Dynamo .
Destination DB Engine	Select GeminiDB Cassandra .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select an available GeminiDB Cassandra instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <p>Full: All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-388 AZ




Table 3-474 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-389 Enterprise Project and Tags




Table 3-475 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-390 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

region

accessKey

secretAccessKey

This button is available only after the replication instance is created successfully.

Table 3-476 Source database settings

Parameter	Description
IP Address or Domain Name	The domain name of the source database. The value is in the format of dynamodb.<region>.amazonaws.com . For example, dynamodb.ap-northeast-1.amazonaws.com .
Port	The port number of the source database. The default value is 80 .
region	The region where the source database resides, for example, ap-northeast-3 .

Parameter	Description
accessKey	IAM key for accessing the source DynamoDB database.
secretAccessKey	

Figure 3-391 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 3-477 Destination database settings

Parameter	Description
DB Instance Name	The GeminiDB Cassandra instance selected when you create the synchronization task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The usernames and passwords of the source and destination databases are encrypted and stored in DRS and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-392 Synchronization Mode

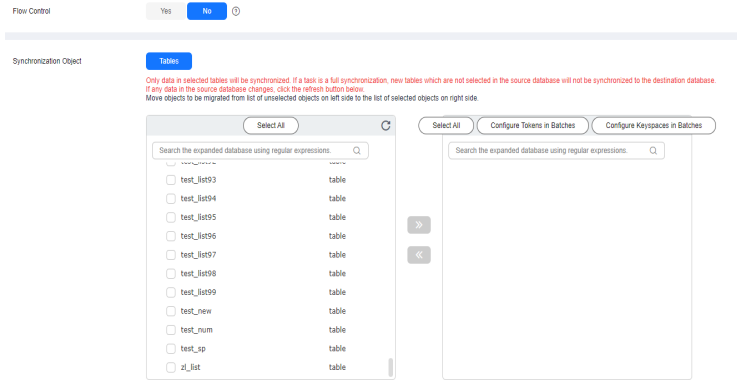
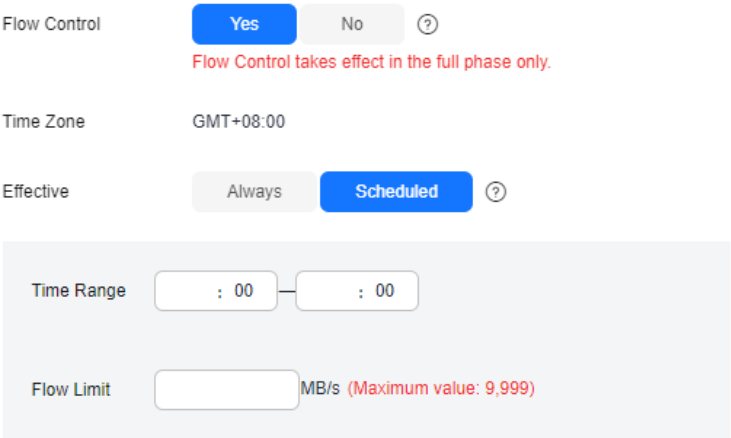


Table 3-478 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-393 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> On the right pane, you can set tokens for selected objects in batches. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). A mapped table name can contain 1 to 48 characters. Only letters, digits, and underscores (_) are allowed. You can select all objects to be synchronized and then set tokens and keyspaces in batches.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

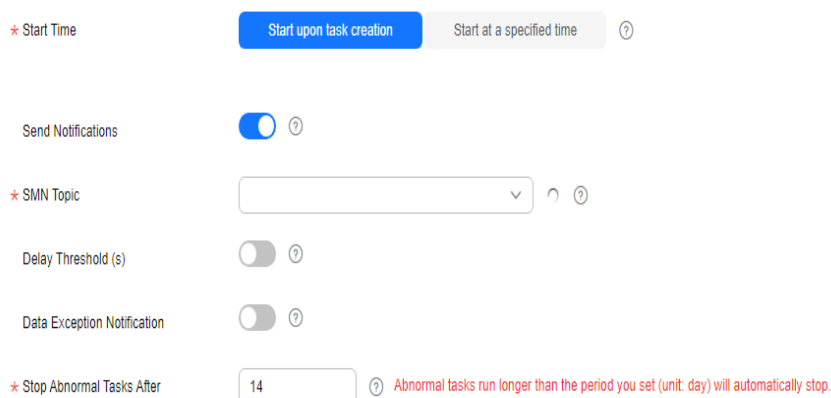
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-394 Task startup settings




The screenshot shows the following settings:

- Start Time:** Two buttons: "Start upon task creation" (active) and "Start at a specified time".
- Send Notifications:** A toggle switch is turned on.
- SMN Topic:** A dropdown menu with a search icon and a help icon.
- Delay Threshold (s):** A toggle switch is turned off.
- Data Exception Notification:** A toggle switch is turned off.
- Stop Abnormal Tasks After:** A text input field containing "14". A red note below it says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 3-479 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.40 From Dynamo to GeminiDB Dynamo

Supported Source and Destination Databases

Table 3-480 Supported databases

Source DB	Destination DB
DynamoDB on other clouds (web services)	GeminiDB Cassandra (DynamoDB API 3.11)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 3-481](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-481 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Supported scenarios: full synchronization and full +incremental synchronization. • Supported fields: String, Number, Binary, Boolean, Null, Map, List, String Set, Number Set, and Binary Set. • Table-level synchronization is supported. <ul style="list-style-type: none"> – During full synchronization, only table data and table structures can be synchronized. – TTL synchronization is not supported. – Counter tables cannot be synchronized.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-482](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-482 Database user permissions

Type	Full and Full+Incremental
Source database user	AmazonDynamoDBFullAccess grants access to all tables and features. You can also create a custom policy to grant only access to required resources. The user must have the following minimum permissions: DescribeTable, Scan, GetItem, DescribeStream, GetRecords, ListStreams, GetShardIterator, ListTables and DescribeEndpoints.
Destination database user	By default, GeminiDB Cassandra (DynamoDB API) does not require authentication. To enable authentication, contact GeminiDB Cassandra technical support.

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-483 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> • Considering the overall synchronization speed and the validity period of the source DynamoDB Stream data, you need to select a proper number and size of tables for concurrent synchronization. It is recommended that the number of tables to be synchronized be no more than 50. • Table structure synchronization includes synchronization of partition keys, sort keys, global secondary indexes, and local secondary indexes. • If full+incremental data synchronization is required, start the streams of the tables to be synchronized in the source database in advance and ensure that the information written to the streams is the old and new images of the project.

Type	Restrictions
Full synchronization	<ul style="list-style-type: none"> Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. Data read in full synchronization uses strong consistency.
Incremental synchronization	<ul style="list-style-type: none"> Do not change the port of the destination database, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. Do not write data to the destination database. Otherwise, data may be inconsistent. To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Procedure

This section uses real-time synchronization from DynamoDB to GeminiDB Cassandra (DynamoDB API) as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 3-395 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-484 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 3-396 Synchronization instance information

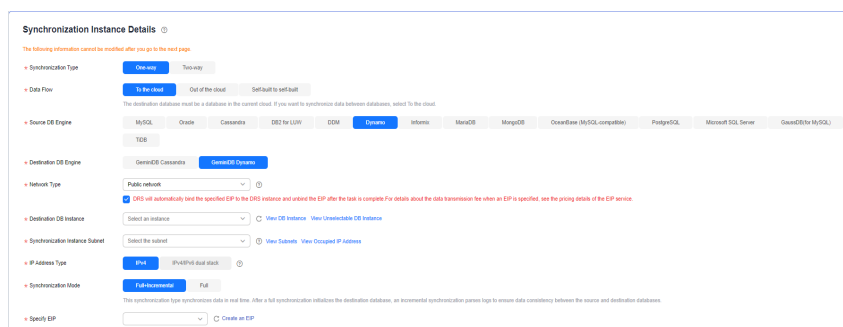


Table 3-485 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Dynamo .
Destination DB Engine	Select GeminiDB Dynamo .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select an available GeminiDB Cassandra instance. GeminiDB Cassandra supports the DynamoDB API.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Full: All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. Full+Incremental: This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-397 AZ



Table 3-486 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-398 Enterprise Project and Tags




Table 3-487 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-399 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

region

accessKey

secretAccessKey

Table 3-488 Source database settings

Parameter	Description
IP Address or Domain Name	The domain name of the source database. The value is in the format of dynamodb.<region>.amazonaws.com . For example, dynamodb.ap-northeast-3.amazonaws.com .
Port	The port number of the source database. The default value is 80 .
region	The region where the source database resides, for example, ap-northeast-3 .

Parameter	Description
accessKey	IAM key for accessing the source DynamoDB database.
secretAccessKey	

Figure 3-400 Destination database information

Destination Database

DB Instance Name

accessKey

secretAccessKey

This button is available only after the replication instance is created successfully.

Table 3-489 Destination database settings

Parameter	Description
DB Instance Name	The GeminiDB Cassandra instance selected when you create the synchronization task. This parameter cannot be changed.
accessKey	IAM key for accessing the destination DynamoDB database. By default, GeminiDB Cassandra (DynamoDB API) does not require authentication. To enable authentication, contact GeminiDB Cassandra technical support.
secretAccessKey	

NOTE

The usernames and passwords of the source and destination databases are encrypted and stored in DRS and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Figure 3-401 Synchronization Mode

Flow Control Yes No ⓘ

Synchronization Object

If you rename tables in the source database, the tables cannot be synchronized. You need to edit this synchronization task to select the renamed tables. If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All ⓘ

Search the expanded database using regular expressions.

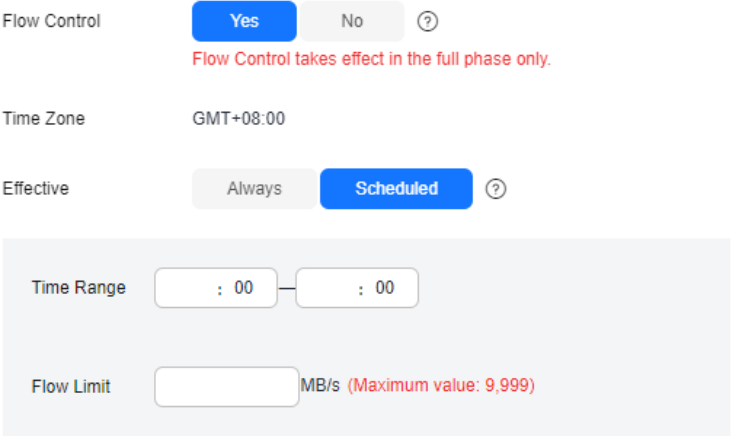
- 0703_music0 table
- MyTable-0_9 table
- MyTable-0_9My_Table_0_9My_Table_0... table
- test_blob table
- test_list table
- test_num table
- test_sp table

Select All ⓘ

Search the expanded database using regular expressions.

- Music [Edit](#) table
- sz_movies0 table
- z_list table

Table 3-490 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-402 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). A mapped table name can contain 1 to 48 characters. Only letters, digits, and underscores (_) are allowed.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

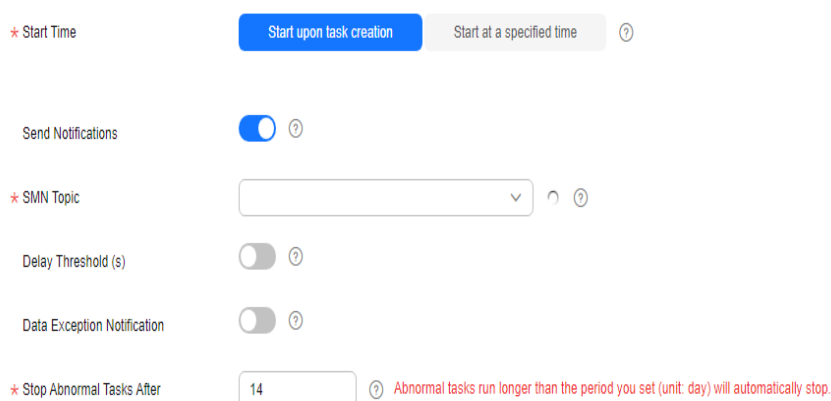
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-403 Task startup settings




The screenshot shows the following settings:

- * Start Time:** Two radio buttons are present: "Start upon task creation" (selected) and "Start at a specified time".
- Send Notifications:** A toggle switch is turned on.
- * SMN Topic:** A dropdown menu is shown with a downward arrow and a refresh icon.
- Delay Threshold (s):** A toggle switch is turned off.
- Data Exception Notification:** A toggle switch is turned off.
- * Stop Abnormal Tasks After:** A text input field contains the value "14". A tooltip below it reads: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 3-491 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none">• If the delay threshold is set to 0, no notifications will be sent to the recipient.• In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.• Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none">• You can set this parameter only for pay-per-use tasks.• Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.41 From AWS DocumentDB to DDS

Supported Source and Destination Databases

Table 3-492 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • AWS DocumentDB (Replica set 4.0, replica set 5.0 and cluster 5.0) Note: <ol style="list-style-type: none"> 1. The default version of an AWS DocumentDB cluster is 5.0. 2. Replica sets support full+incremental, full, and incremental synchronization modes. 3. Clusters support only the full synchronization mode. 4. An Amazon DocumentDB instance-based cluster corresponds to a DDS replica set. 5. An Amazon DocumentDB elastic cluster corresponds to a DDS cluster. 	DDS DB instances (versions 3.4, 4.0, 4.2, and 4.4) NOTE The destination database version must be the same as or later than the source database version.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-493](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-493 Database account permissions

Type	Full+Incremental
Source database user	<p>Replica set: The source database user must have the readAnyDatabase permission for the admin database.</p> <p>Cluster: The source database user must have the readAnyDatabase permission for the admin database.</p>
Destination database user	The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 3-494](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-494 Supported synchronization objects

Type	Notes
Synchronization objects	<ul style="list-style-type: none"> • Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. • System databases (such as local, admin, and config) and system collections cannot be synchronized. If service data is in a system database, run renameCollection to move the service data to the user database. • The statement for creating a view cannot contain a regular expression. • Collections that contain the _id field without indexes are not supported. • The first parameter of BinData() cannot be 2. • Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see "Pre-check Items". In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-495](#).

Table 3-495 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • During incremental synchronization, change streams of the source database must be enabled. <ol style="list-style-type: none"> 1. To enable change streams for all collections on a specific database, authenticate to your Amazon DocumentDB cluster using the mongo shell and execute the following command: (replace <i>db_name</i> with the actual database name): <code>db.adminCommand({modifyChangeStreams: 1,database: "db_name",collection: "", enable: true});</code> If <i>collection</i> is set to an empty string, change streams is enabled for all collections within a given database. To enable change streams for a specific collection, specify the collection name. For more information, click here. 2. The change streams retention period can be changed. The default retention period is three hours. You need to evaluate the migration duration based on the data volume and change the retention period to a value greater than or equal to the migration duration. If you do not know the time required for the migration, you are advised to set the retention period to seven days. After the migration is complete, change the retention period to an appropriate value. You can change the retention period through the console or CLI. For details, click here.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database, such as deleting databases, collections, indexes, documents, or views. Otherwise, the synchronization may fail. ● During the synchronization, data rollback caused by a primary/standby switchover of the source database is not supported. ● Documents larger than 16 MB in the source database cannot be inserted or updated during full or incremental synchronization. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● If the source is a replica set instance, the following operations and commands are supported during incremental synchronization: <ul style="list-style-type: none"> - Adding, deleting, and updating documents - Deleting collections - The dropDatabase, dropCollection, and renameCollection DDL commands ● In the incremental synchronization phase, concurrent replay is performed at the collection level to maintain the synchronization performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. ● In either of the preceding scenarios, the task delay may increase. <p>Troubleshooting</p> <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● Replica set: The AWS DocumentDB replica set instance must be available and have primary nodes. ● The value of block_compressor is determined by stats().wiredTiger.creationString.block_compressor of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be synchronized. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on net.compression.compressors of the destination database. If the destination database version is DDS 4.2, DRS does not synchronize compression parameters because the destination database does not support compression parameter settings. ● If the AWS DocumentDB service of the source database is deployed with other services on the same server, set the value of the cacheSizeGB parameter to the half of the minimum idle cache for the WiredTiger engine of the source database. ● During an incremental synchronization of collections, renaming collections is not recommended. ● Specify the information about all primary and secondary nodes of the replica set to prevent the primary/standby switchover from affecting the synchronization task. If you enter information about multiple primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● To accelerate the synchronization, delete unnecessary indexes from the source database and retain only necessary indexes before the synchronization. You are advised not to create indexes for the source database during the synchronization. If indexes must be created, create them in the background. ● To prevent loopback, do not start tasks that migrate the same database to and out of the cloud at the same time. <p>NOTE The source and destination databases in a synchronization task can be the same database. To avoid replication loop, you must rename object names.</p>

Procedure

This section uses DDS replica set as an example to describe how to configure an AWS DocumentDB to DDS replica set synchronization task over a VPC network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-404 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 3-496 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 3-405 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

• Synchronization Type: **One-way** Two-way

• Data Flow: **To the cloud** Out of the cloud Self-Built to self-Built
The destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

• Source DB Engine: MySQL Oracle **AWS DocumentDB** Cassandra DB2 for LUW DDM Dynamo Informis MariaDB MongoDB OceanBase (MySQL-compatible) PostgreSQL Microsoft SQL Server GaussDB(for MySQL) TDB

• Destination DB Engine: **DB2**

• Network Type: VPC or Direct Connect

• Destination DB Instance: Select an instance [View DB Instance](#) [View Unselectable DB Instance](#)

• Synchronization Instance Subnet: Select the subnet The IP address is allocated automatically but it can [View Subnets](#) [View Occupied IP Address](#)

• IP Address Type: **IPv4** IPv6/IPv6 dual stack

• Synchronization Mode: **Full/Incremental** Full Incremental
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Source DB Instance Type: **Replica sets**

Table 3-497 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select AWS DocumentDB .
Destination DB Engine	Select DDS .
Network Type	AWS DocumentDB does not support public network access. For details, see Connection issues . You can only select VPN or Direct Connect . <ul style="list-style-type: none"> VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you have created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental , data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.

Parameter	Description
Source DB Instance Type	If Synchronization Mode is set to Full+Incremental , set this parameter based on the source database. Currently, the source database supports only replica set, which is a non-cluster database.

- Task type

Figure 3-406 Task type




Table 3-498 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-407 Enterprise Project and Tags




Table 3-499 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

- Source database information

Figure 3-408 Source database information

Source Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Authentication Database

Database Username

Database Password 🔑

SSL Connection

Table 3-500 Source database settings

Parameter	Description
IP Address or Domain Name	Example: docdb-2024-06-05-03-09-20.cluster-c58eyca06srk.ap-northeast-1.docdb.amazonaws.com:27017
Authentication Database	The name of the authentication database.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-409 Destination database information

Destination Database

DB Instance Name:

Authentication Database:

Database Username:

Database Password:

SSL Connection:

Table 3-501 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the start point and synchronization objects, and click **Next**.

Figure 3-410 Synchronization objects

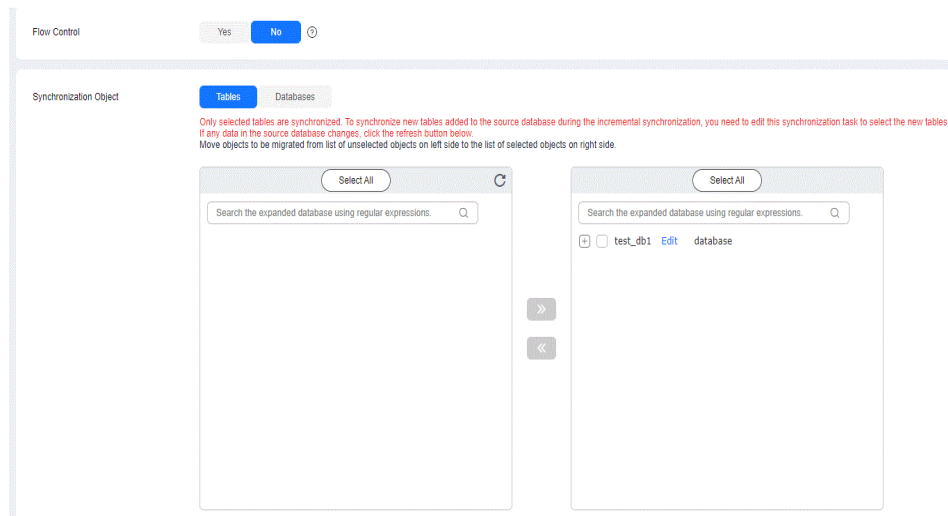
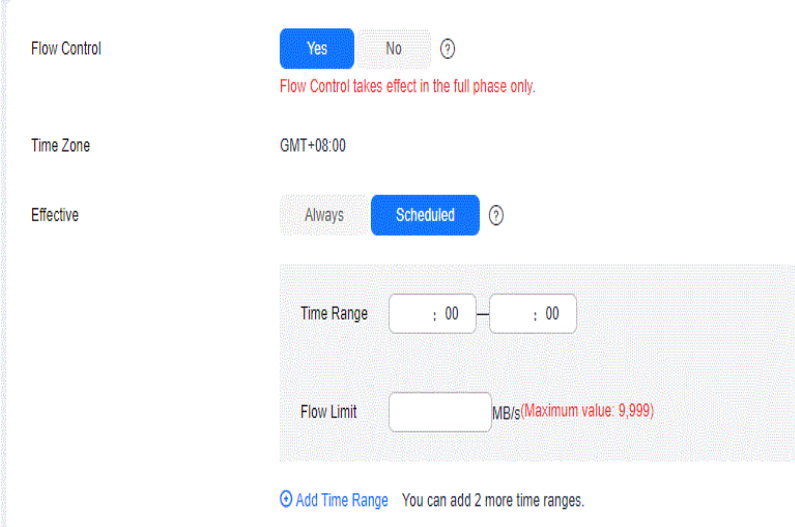




Table 3-502 Synchronization objects

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow control is disabled by default. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> • Yes • You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. • In addition, you can set the time range based on your service requirements. The flow control configuration usually includes settings of a flow control time range and a flow limit value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. • The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-411 Flow control</p>  <ul style="list-style-type: none"> • No • The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which causes read consumption on the source database accordingly. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> • Flow Control takes effect in the full phase only. • You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). The mapped name can contain 1 to 63 characters. The following characters are not allowed \. "\$ < ></p> <p>If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-412 Task startup settings





Table 3-503 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.42 From Redis to GeminiDB Redis

Supported Source and Destination Databases

Table 3-504 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x) ECS-hosted official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x) 	<p>GeminiDB Redis</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p>

 NOTE

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-505 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-505 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Object level: database level • Supported synchronization objects: <ul style="list-style-type: none"> - Data of the String, Hash, List, Set, and Sorted Set types can be synchronized. - Lua scripts can be migrated (except in a single full migration task). - Transactions cannot be synchronized. - Stream, BitMap, HyperLogLog, and GEO data types cannot be synchronized. - User-defined types are not supported. - Commands that are not supported by the destination GeminiDB Redis instance cannot be synchronized.

Suggestions

 CAUTION

To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- During the synchronization, ensure that no data is written to the destination database to keep data consistency before and after the synchronization.
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full synchronization, and two to four CPUs are occupied.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-506 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - To create a full+incremental task, the PSYNC command must be supported. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination must be a GeminiDB Redis instance on the current cloud. - The destination GeminiDB Redis instance must be empty before the migration. - The version of the destination database instance must be the same as or later than that of the source database. ● Other notes: <ul style="list-style-type: none"> - A full migration task reads data in SCAN mode. - A full+incremental task read data by parsing RDB files. - If the source database does not support the PSYNC command, use the full synchronization type. - If the source database supports the PSYNC command, use the full+incremental synchronization type. - The expiration time settings of the full synchronization task cannot be synchronized. - In a full migration task, Lua scripts cannot be migrated. - When synchronizing Lua scripts, ensure that the source database contains other data. Otherwise, Lua scripts in RDB cannot be parsed, causing the Lua script migration failure. - If the source is a self-managed single-node Redis database, the following data types can be migrated during the incremental migration phase: append, blpop, brpop, decr, decrby, del, expire, expireat, flushall, flushdb, getset, hdel, hincrby, hincrbyfloat, hmset, hset, hsetnx, incr, incrby, incrbyfloat, linsert, lpop, lpush, lpushx, lrem, lset, ltrim, move, mset, msetnx, multi, persist, pexpire, pexpireat, psetex, publish, rename, rpop, rpush, rpushx, sadd, select, set, setbit, setex, setnx, setrange, smove, spop, srem, zadd, zincrby, zrem, zremrangebylex, zremrangebyrank and zremrangebyscore. - GeminiDB Redis does not support multi-key commands. If the source is a standalone Redis database and keys involved in an operation aren't in the same hash-slot, the following error message will appear: CROSSSLOT Keys in request don't hash to the same slot - DRS does not perform high-risk operations such as FLUSH on the destination database, so data of the List type may be duplicated.

Type	Restrictions
	<ul style="list-style-type: none"> - To ensure a successful incremental synchronization, increase the value of repl-backlog-size in the redis.conf configuration file of the source database before the synchronization. - During the synchronization, if the self-built Redis database is scaled in or out (for example, shards are added or deleted) or the specifications are changed (for example, the memory is scaled up), reconfigure the synchronization task. To ensure data consistency, clear the data that has been migrated to the destination GeminiDB Redis before reconfiguring the task. - If the expire policy is configured for some keys in the source database, the number of keys in the destination database (for example, by running the info command) may be less than that in the source database because the keys may have expired but are not deleted. To ensure data consistency, you are advised to prolong the key expiration time.
Full migration	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During the synchronization, do not modify the destination database.
Incremental migration	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During the synchronization, do not modify the destination database.

Prerequisites

- [You have logged in to the DRS console.](#)
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management](#).

Procedure

This section describes how to use DRS to synchronize data from a Redis database to GeminiDB Redis over a public network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 3-413 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 3-507 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\\"

- Synchronization instance information

Figure 3-414 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you click the next page.

Synchronization Type:

Data Flow: To the cloud Out of the cloud Self-built to self-built

The destination database must be a database in the current cloud. If you want to synchronize data between databases, select To the cloud.

Source DB Engine:

Destination DB Engine:

Network Type:

Destination DB Instance:

Synchronization Instance Name:

IP Address Type: ⓘ

Synchronization Mode:

This synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization starts to ensure data consistency between the source and destination databases.

Specify EP:

Table 3-508 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select Single-Node or Master/Standby Redis .
Destination DB Engine	Select GeminiDB Redis .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Full: All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. <p>NOTE If you are performing a full synchronization, do not perform operations on the source database. Otherwise, data generated in the source database during the synchronization will not be synchronized to the destination database.</p> - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-415 AZ

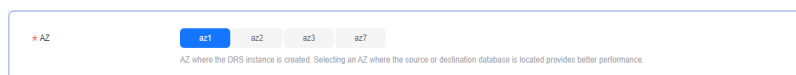


Table 3-509 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-416 Enterprise Project and Tags




Table 3-510 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Configuring your own DNS server

Figure 3-417 DNS Server

Configure Your Own DNS Server ⓘ

DNS Server

DNS Server IP Address . .

Table 3-511 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use your own private domain name as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data synchronization.

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 3-418 Source database information

Source Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Password ⓘ

Table 3-512 Source database information

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Password	Password used to log in to the source database.

 NOTE

The IP address or domain name, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Destination database configuration

Figure 3-419 Destination database information

Destination Database

DB Instance Name

Database Password

Table 3-513 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Database Password	The password for the database username.

 NOTE

The password of the destination database is encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the password are permanently deleted.

- Step 4** On the **Set Synchronization Task** page, configure the objects to be synchronized, and then click **Next**.

Figure 3-420 Synchronization objects

Synchronization Object


Newly added tables in the database will be automatically synchronized to the destination, and deleted tables are automatically no longer synchronized.
If any data in the source database changes, click the refresh button below.
Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Search the expanded database using regular expressions.

- 0 database
- 1 database
- 10 database
- 11 database
- 12 database
- 13 database
- 14 database
- 15 database
- 2 database
- 3 database
- 4 database
- 5 database

Search the expanded database using regular expressions.

Table 3-514 Synchronization Object

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports database-level synchronization. You can select the databases to be synchronized based on service requirements.</p> <p>If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, and click **Submit** to submit the task.

Figure 3-421 Task startup settings



Start Time: Start upon task creation Start at a specified time ⓘ

Apr 11, 2024 03:00:00 ⓘ GMT+08:00

In the full migration phase, the DRS task occupies performance resources. To ensure services and tasks can run normally, you are advised to start the task during off-peak hours and reserve two to three days for data verification.

Send Notifications: ⓘ

* SMN Topic: ⓘ ⓘ

Delay Threshold (s): ⓘ


Data Exception Notification:

* Stop Abnormal Tasks After: ⓘ Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 3-515 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements. The Start at a specified time option is recommended.</p> <p>NOTE The synchronization task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the task billing is about to start, the status, latency metric, or data of the synchronization task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. • If the delay threshold is set to 0, no notifications will be sent to the recipient.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.</p>

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.43 From Redis Cluster to GeminiDB Redis

Supported Source and Destination Databases

Table 3-516 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises open-source Redis Cluster (versions 3.0.x, 4.0.x, and 5.0.x) • ECS-hosted open-source Redis Cluster (versions 3.0.x, 4.0.x, and 5.0.x) 	GeminiDB Redis NOTE The destination database version must be the same as or later than the source database version.

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-517 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-517 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Object level: All• Supported synchronization objects:<ul style="list-style-type: none">- Data of the String, Hash, List, Set, and Sorted Set types can be migrated.- Lua scripts can be migrated.- Transactions cannot be migrated.- Data of the Stream, BitMap, HyperLogLog, and GEO types cannot be migrated.- User-defined types are not supported.- Commands that are not supported by the destination GeminiDB Redis instance cannot be migrated.

Suggestions

CAUTION

To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- During the synchronization, ensure that no data is written to the destination database to keep data consistency before and after the synchronization.
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full synchronization, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single

full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-518 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - To create a full+incremental task, the PSYNC command must be supported. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination must be a GeminiDB Redis instance on the current cloud. - It is recommended that the destination GeminiDB Redis instance be empty before the synchronization. - The version of the destination database instance must be the same as or later than that of the source database. ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files. - When migrating Lua scripts, ensure that the source database contains other data. Otherwise, Lua scripts in RDB cannot be parsed, causing the Lua script synchronization failure. - DRS does not perform high-risk operations such as FLUSH on the destination database, so data of the List type may be duplicated. - When testing the connection to the source database, you must enter the IP addresses of all shards in the Redis Cluster. To keep services accessible to the source Redis instance, you are advised to synchronize data from the replica node (slave node) of the source instance. That is, you are advised to enter the IP address of the slave node in the Redis Cluster instance. - To ensure a successful incremental synchronization, increase the value of repl-backlog-size in the redis.conf configuration file of the source database before the synchronization. - During the synchronization, if the source Redis Cluster instance is scaled in or out (for example, shards are added or deleted) or the specifications are changed (for example, the memory is scaled up), reconfigure the synchronization task. To ensure data consistency, clear the data that has been migrated to the destination GeminiDB Redis before reconfiguring the task. - If the expire policy is configured for some keys in the source database, the number of keys in the destination database (for example, by running the info command) may be less than that in the source database because the keys may have expired but are not deleted. To ensure data consistency, you are advised to prolong the key expiration time.

Type	Restrictions
Full migration	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During synchronization, if a primary/standby switchover occurs in the source Redis Cluster instance, the task will fail. In this case, you need to clear the destination database to reconfigure the task. • During the synchronization, do not modify the destination database.
Incremental migration	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During synchronization, if a primary/standby switchover occurs in the source Redis Cluster instance, the task will fail. In this case, you need to clear the destination database to reconfigure the task. • During the synchronization, do not modify the destination database.

Prerequisites

- [You have logged in to the DRS console.](#)
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management](#).

Procedure

This section describes how to use DRS to synchronize data from a Redis Cluster instance to GeminiDB Redis over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 3-422 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-5678

Description: [Text area]

0/256

Table 3-519 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

- Synchronization instance information

Figure 3-423 Synchronization instance information

Synchronization Instance Details

⚠ The following information cannot be modified after you go to the next page.

Synchronization Type: **On-prem** | **On-cloud** | **Hybrid**

Data Flow: **To the cloud** | Out of the cloud | Self-host to self-built

Source DB Engine: MySQL, Oracle, Cassandra, DCS for LUPP, EDW, Dynamo, Informatica, MariaDB, MongoDB, OceanBase (MySQL compatible), PostgreSQL, Single-Node or Master/Standby Redis, **ClickHouse**, Microsoft SQL Server, GaussDB (MySQL), TDS

Destination DB Engine: **ClickHouse**

Network Type: Public network

Destination DB Instance: Select an instance | View DB Instance | View Unavailable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnet

IP Address Type: **IPv4** | IPv6

Synchronization Mode: **Full Synchronization**

DRS Task Nodes: 2

Task EP: **Create an EP**

Table 3-520 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select Redis cluster .
Destination DB Engine	Select GeminiDB Redis .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>

Parameter	Description
Source Shard Quantity	Number of subtasks connected to the source Redis Cluster instance. The value ranges from 1 to 16 and cannot be greater than the number of shards in the source Redis Cluster instance. Set this parameter based on the scale of the source Redis Cluster instance. You are advised to set one subtask to connect to four shards in the source Redis Cluster instance.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 3-424 AZ




Table 3-521 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 3-425 Enterprise Project and Tags




Table 3-522 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 3-426 DNS Server




Table 3-523 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use the IP address of your own DNS server as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data migration.

 **NOTE**

This function is available when you need to use the IP address of your own DNS server as the source or destination database IP address.

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 3-427 Source database information

Source Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Password ⓘ

Table 3-524 Source database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the source database in the IP address/Domain name:Port format. Enter the IP addresses and ports of all shards in the source Redis Cluster instance. You are advised to enter the IP addresses of all shard slave nodes. Up to 32 IP addresses or domain names can be entered. Use commas (,) to separate multiple IP addresses or domain names. For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Password	Password used to log in to the source database.

 **NOTE**

The IP address or domain name, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Destination database configuration

Figure 3-428 Destination database information

Destination Database

DB Instance Name

Database Password

Table 3-525 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Database Password	The password for the database username.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the password are permanently deleted.

- Step 4** On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 3-429 Synchronization objects

Synchronization Object

Table 3-526 Synchronization Object

Parameter	Description
Synchronization Object	All database objects can be synchronized.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, and click **Submit** to submit the task.

Figure 3-430 Task startup settings




Table 3-527 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. The Start at a specified time option is recommended. NOTE The synchronization task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the task billing is about to start, the status, latency metric, or data of the synchronization task is abnormal, DRS will send you a notification.

Parameter	Description
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. If the delay threshold is set to 0, no notifications will be sent to the recipient.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4 Out of the Cloud

4.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 4-1 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">RDS for MySQL (5.5, 5.6, 5.7, and 8.0)	<ul style="list-style-type: none">On-premises MySQL (5.5, 5.6, 5.7, and 8.0)ECS-hosted MySQL (5.5, 5.6, 5.7, and 8.0)Other cloud MySQL (5.5, 5.6, 5.7, and 8.0)RDS for MySQL (5.5, 5.6, 5.7, and 8.0)

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-level comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.
- If you create a many-to-one synchronization task, refer to [Constraints and Operation Suggestions on Many-to-One Scenario](#).
- For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed. Otherwise, all synchronization tasks fail.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-2 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"><li data-bbox="587 344 1369 472">• The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT<li data-bbox="587 488 1414 723">• The destination database user must have the following permissions: SELECT, CREATE, INDEX, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, primary key indexes, unique indexes, common indexes, store procedures, views, and functions can be synchronized. ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. ● Only MyISAM and InnoDB tables can be synchronized. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. ● Braces ({}) in a database or table name in the source database must be used in pairs. Otherwise, the table structure synchronization may fail. After data mapping, if braces ({}) in a database or name mapped to the destination database are not in pairs, the table structure may fail to be synchronized. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database names cannot contain non-ASCII characters, or the following characters: <code>.'<>\/\</code>" ● The source table and view names cannot contain non-ASCII characters, or the following characters: <code>.'<>\/\</code>" ● The column names in the source database tables cannot end with a backslash (<code>\</code>). ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During the synchronization, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The source database cannot be a read replica.

Type	Restrictions
Destination database	<ul style="list-style-type: none"> • Data cannot be synchronized from a newer version database to an older version database. • The destination database must have sufficient disk space. • The character set of the destination database must be the same as that of the source database. • The time zone of the destination database must be the same as that of the source database. • If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. • During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. • If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (')

Type	Restrictions
Precautions	<ul style="list-style-type: none">• If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created.• If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see Syntax Differences Between MySQL Versions.• Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys• If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent.• If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported.• If the sources and destinations are RDS instances, database mapping is required.• The source and destination databases cannot contain tables that have the same names but do not have primary keys.• If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization.• If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem.• If the destination MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.• If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss.• Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail.• Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key.

Type	Restrictions
	<ul style="list-style-type: none"> ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● The source database does not support point-in-time recovery (PITR). ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● During the full and incremental synchronization phases, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. ● After a task is created, the destination database cannot be set to read-only. ● A real-time synchronization task may fail due to the change of the username and password of the source or destination database. If it happens, rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization. ● If the source or destination database port is changed during data synchronization, the synchronization task fails. You can rectify the fault as follows:

Type	Restrictions
	<ul style="list-style-type: none"> - If the source database port is wrong, correct the port number on the DRS console and then retry the synchronization task. - If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization. • To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during synchronization. • DDL operations are not supported during full synchronization. • During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • You can add additional objects during an incremental synchronization. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. • During database name mapping, function objects, views, and stored procedures may fail to be created because they are referenced by a database table name before the mapping. DRS ignores the error reported during the creation of these objects in the full synchronization phase. As a result, these objects are lost in the destination database. In the incremental synchronization phase, if database name mapping is required, DDL operations (including CREATE, MODIFY, and DELETE) related to function objects, views, and stored procedures are not synchronized to the destination database.

Procedure

This section uses synchronization from RDS for MySQL to MySQL as an example to describe how to use DRS to create a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-1 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region ..
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project ..

* Task Name ⓘ

Description
0/256

Table 4-3 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-2 Synchronization instance details

Synchronization Instance Details ⓘ

⚠ The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud Self-built to self-built
Out of the cloud: The source database must be a database in the current cloud.

• Source DB Engine MySQL ODM GaussDB Distributed GaussDB Primary/Standby MariaDB DDS PostgreSQL GaussDB for MySQL

• Destination DB Engine MySQL Oracle CSSES Kafka MariaDB

• Network Type ⓘ
 DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

• Source DB Instance ⓘ [View DB Instance](#) [View Unavailable DB Instance](#)

• Synchronization Instance Subnet ⓘ [View Subnets](#) [View Occupied IP Address](#)

• Security Group ⓘ

• Synchronization Mode Full-synchronous Incremental
The synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization period logs to ensure data consistency between the source and destination databases.

• Specify EP ⓘ [Create an EP](#)

Table 4-4 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-3 Task type




Table 4-5 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-4 Enterprise Project and Tags




Table 4-6 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

In different data flow scenarios, the source and destination database settings are different. Specify the required parameters based on the GUI.

- Source database information

Figure 4-5 Source database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 4-7 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

- Destination database information

Figure 4-6 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 4-8 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the destination database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 4-7 Synchronization mode

Flow Control Yes No ⊙

Incremental Conflict Policy Ignore Overwrite ⊙
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Filter DROP DATABASE Yes No

Synchronize Normal index Incremental DDLs ⊙

Create Indexes Along With Table Structure Yes No

Synchronize DML Insert Update Delete ⊙

Online DDL Yes No
Online DDL can be configured during table-level synchronization.

Data Synchronization Topology One-to-one One-to-many Many-to-one ⊙

Synchronize DDLs Default Custom ⊙
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. The following shows only part of DDL statements.

<input checked="" type="checkbox"/> Add	<input checked="" type="checkbox"/> CREATE_TABLE <input checked="" type="checkbox"/> ADD_PRIMARY_KEY	<input checked="" type="checkbox"/> ADD_COLUMN	<input checked="" type="checkbox"/> ADD_INDEX	<input checked="" type="checkbox"/> CREATE_INDEX
<input checked="" type="checkbox"/> Modify	<input checked="" type="checkbox"/> RENAME_TABLE <input checked="" type="checkbox"/> RENAME_COLUMN	<input checked="" type="checkbox"/> MODIFY_COLUMN	<input checked="" type="checkbox"/> CHANGE_COLUMN	<input checked="" type="checkbox"/> RENAME_INDEX
<input checked="" type="checkbox"/> Delete (High-risk)	<input checked="" type="checkbox"/> DROP_COLUMN <input checked="" type="checkbox"/> DROP_PRIMARY_KEY	<input checked="" type="checkbox"/> DROP_INDEX <input checked="" type="checkbox"/> DROP_PARTITION	<input checked="" type="checkbox"/> DROP_TABLE	<input checked="" type="checkbox"/> TRUNCATE_TABLE

Synchronization Object Tables Databases Import object file
If any data in the source database changes, click the refresh button below.
 Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All ⊙

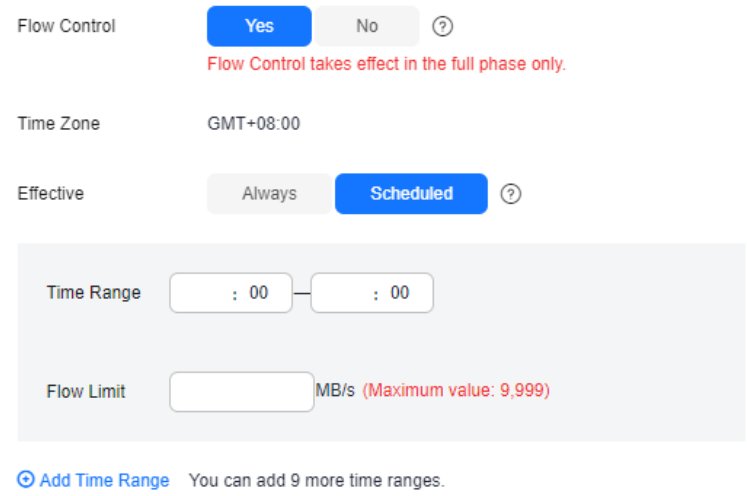
Search the expanded database using regular expressions.

test01 database

Select All ⊙

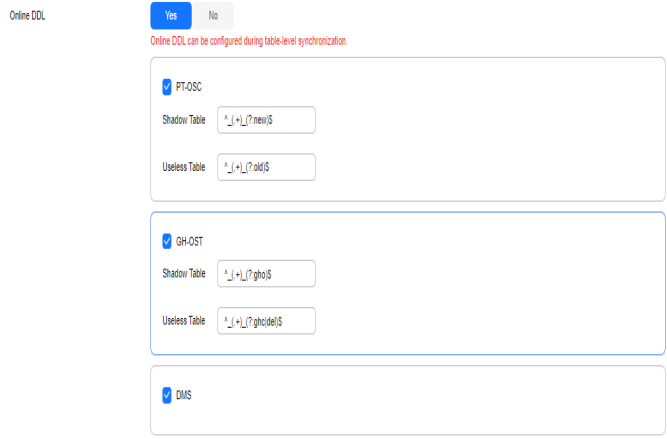
Search the expanded database using regular expressions.


Table 4-9 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-8 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none"> ● Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase. ● No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OSC, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 4-9 Online DDL</p>  <p>The screenshot shows the 'Online DDL' configuration page. At the top, there are 'Yes' and 'No' buttons, with 'Yes' selected. Below this is a red note: 'Online DDL can be configured during table-level synchronization'. There are three tool configuration sections:</p> <ul style="list-style-type: none"> PT-OSC: Checked. Shadow Table: <code>^_+_?(?mev)\$</code>. Useless Table: <code>^_+_?(?old)\$</code>. GH-OST: Checked. Shadow Table: <code>^_+_?(?gho)\$</code>. Useless Table: <code>^_+_?(?ghode)\$</code>. DMS: Checked. No: Table-level synchronization does not support Online DDL synchronization.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-10 Processing Data



Belonged Database/Table	New Name	Column Name	Operation Type	Type	Operation
gtest.order_payment_1	gtest.order_payment_1	c1	The serverName@database@table column is used.	varchar(191)	Add Delete
gtest.order_payment_10	gtest.order_payment_10	c2	Default	int	Add Delete
gtest.order_payment_11	gtest.order_payment_11	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-11 Task startup settings

Table 4-10 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.2 From MySQL to Kafka

Supported Source and Destination Databases

Table 4-11 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">RDS for MySQL (5.5, 5.6, 5.7, and 8.0)	<ul style="list-style-type: none">Kafka

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-12 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, SHOW VIEW, EVENT, REPLICATION SLAVE, and REPLICATION CLIENT.
Synchronization object	<ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The database and table names in the source database cannot contain non-ASCII characters, or the following special characters: <code>!<>\/</code> ● The column names in the source database tables cannot end with a backslash (<code>\</code>). ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.
Destination database	<ul style="list-style-type: none"> ● The destination database is a Kafka database. ● You are advised to set auto.create.topics.enable of Kafka to false.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the data types are incompatible, the synchronization may fail. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● A real-time synchronization task may fail due to the change of the username and password of the source database. You need to rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization. ● If the source database port is changed during data synchronization, the synchronization task fails. If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization. ● If a real-time synchronization task fails as the IP address is changed, the system automatically changes the IP address to the correct one. Then, you need to retry the task to continue the synchronization. Therefore, changing the IP address is not recommended. ● If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database does not support point-in-time recovery (PITR). ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● Data inconsistency may occur when the MyISAM table is modified during synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> • During synchronization of table-level objects, renaming tables is not recommended. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-12 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ 0/256

Table 4-13 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-13 Synchronization instance details

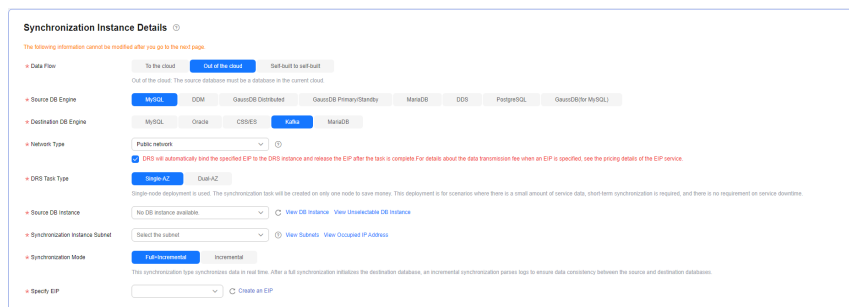


Table 4-14 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select Kafka .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 4-14 Task type





Table 4-15 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-15 AZ</p> 

- Enterprise Project and Tags

Figure 4-16 Enterprise Project and Tags




Table 4-16 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-17 Source database information

Source Database

DB Instance Name

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-17 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-18 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

● Test successful

Table 4-18 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

The parameters on the **Set Synchronization Task** page vary according to the synchronization mode selected in **Step 2**. The **Full+Incremental** synchronization mode is used as an example in **Figure 4-19**.

Figure 4-19 Synchronization mode

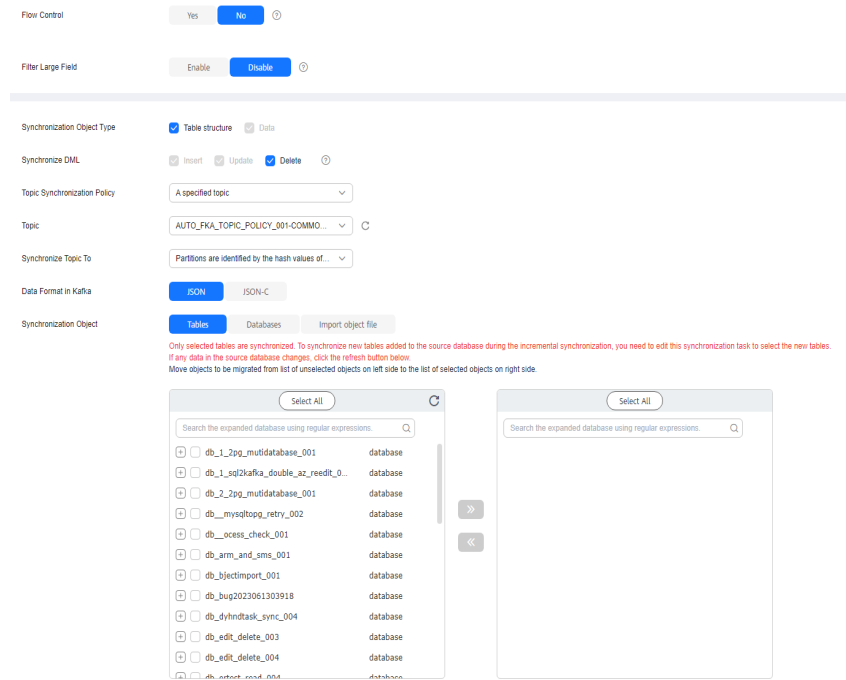
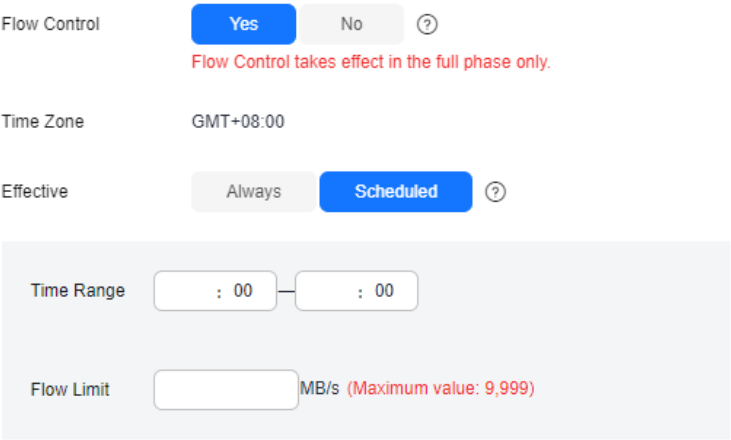
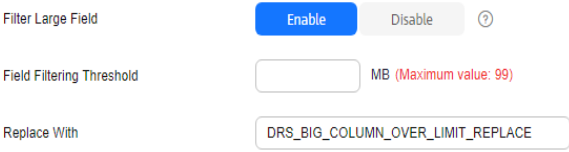



Table 4-19 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-20 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Filter Large Field	<p>Indicates whether to use large field filtering to process special fields (blob, mediumblob, longblob, varbinary, mediumtext and longtext) in a synchronization table.</p> <ul style="list-style-type: none"> Enable You need to set Field Filtering Threshold and Replace With. If the size of a field exceeds the threshold, the value is replaced based on a specified character. Note that large field filtering is used to replace the value of a field, not the entire DML record. If a DML record contains many large fields, the size of only some of these fields exceeds the filtering threshold, and the accumulated value of other fields that do not exceed the filtering threshold is greater than the value of request.max.size, when data is written to Kafka, the size of the message body in the destination Kafka database may still exceed the upper limit, resulting in a DRS error. Figure 4-21 Setting large field filtering  Disable Large fields are not filtered.
Synchronization Object Type	<p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name.</p> <p>For example, if this parameter is set to \$database\$-tablename\$ and the database name is db1, and the table name is tab1, the topic name is db1-tab1. If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>


Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • Partitions are identified by the hash values of the primary key: This mode applies to scenarios where a single table contains a large amount of data, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. Data sequence can be preserved only when the primary key value is not changed. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.table_name.
Data Format in Kafka	<p>Select the data format to be delivered from MySQL to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <ul style="list-style-type: none"> When importing an object file, you can perform either topic mapping or object name mapping. If you perform topic mapping when importing an object file, different tables can be synchronized to different topics in the destination database. If topic mapping is not specified for an object, the object uses the external topic policy. You can modify the mapping when editing the synchronization object. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 4-22 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-23 Task startup settings

Table 4-20 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.3 From MySQL to CSS/ES

Supported Source and Destination Databases

Table 4-21 Supported databases

Source DB	Destination DB
RDS for MySQL (5.5, 5.6, 5.7, and 8.0)	ElasticSearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9 and 7.10

 NOTE

Only whitelisted users can use this function.

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-22 Precautions

Type	Restrictions
Database permissions	Minimum permission requirements for full plus incremental synchronization: <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the following permissions: READ and WRITE

Type	Restrictions
Synchronization object	<ul style="list-style-type: none">• The table data can be synchronized.• Databases, views, indexes, constraints, functions, stored procedures, triggers, and events cannot be synchronized.• The system database and event status cannot be synchronized.• Tables whose primary keys are of the FLOAT type cannot be synchronized.• Tables that do not have primary keys cannot be synchronized.• Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>/' ● The column name in the source database cannot contain non-ASCII characters or the following characters: '.' ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails.

Type	Restrictions
	<ul style="list-style-type: none">• The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.• The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail.• The source database cannot be a read replica.
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly.• The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored to a point in time when a full synchronization was being performed. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Convert the value range of the source database time field that is not supported by the destination database to null. ● The strings padded with zeros in the source database may be truncated because the source database uses the fixed-length binary data type, and the destination database uses the variable-length data type. ● The binary value is encrypted using Base64 and then written to the destination database. ● If no time zone is specified for the source database, specify the time zone for synchronizing the datetime type to the destination database. ● All table field names are converted to lowercase letters. ● If the <code>_id</code> field of the destination database is generated using multiple columns in the source database, separate these columns with colons (:). ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● During the synchronization, do not create a database named ib_logfile in the source. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.

Type	Restrictions
	<ul style="list-style-type: none"> • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable upload is supported, but data may be repeatedly inserted into non-transactional tables that do not have primary keys when the server system breaks down. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-24 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-23 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-25 Synchronization instance details

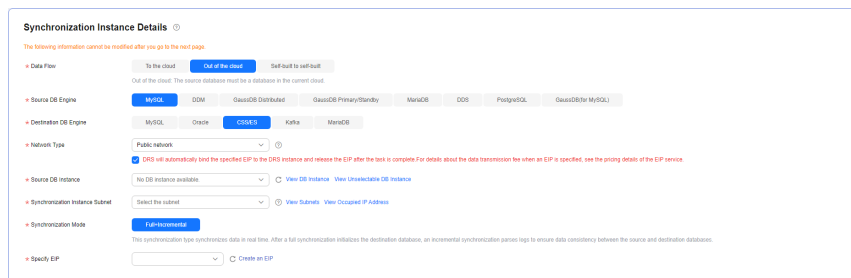


Table 4-24 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select CSS/ES .

Parameter	Description
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-26 Task type



Table 4-25 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-27 Enterprise Project and Tags




Table 4-26 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-28 Source database information

Source Database

DB Instance Name

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-27 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-29 Destination database information

Destination Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-28 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. Only .cer and .pem certificates are supported. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-30 Synchronization mode

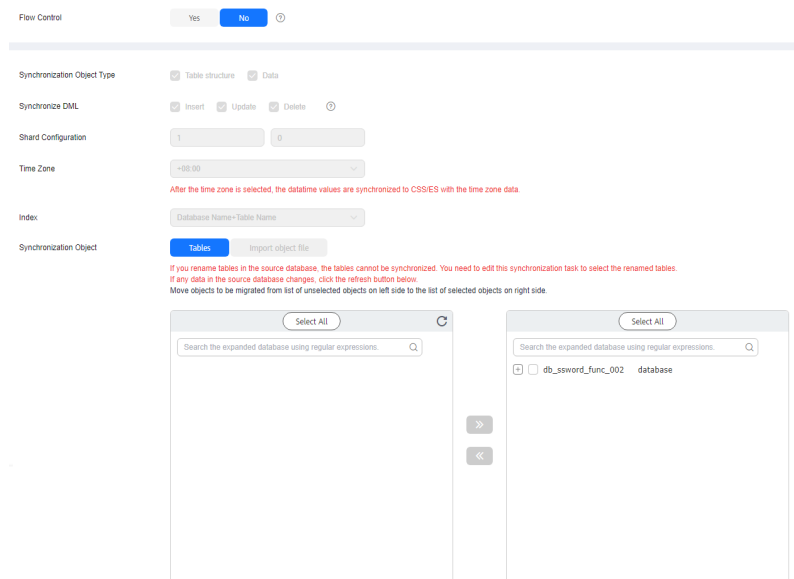
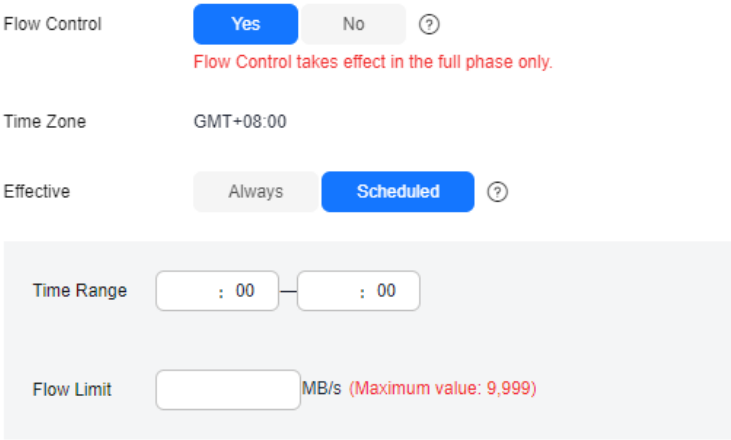



Table 4-29 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-31 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.


Parameter	Description
Synchronization Object Type	<p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Shard Configuration	<p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p>
Time Zone	<p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p>
Index Name	<ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> You can select objects by importing a file. For details, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). If you select Tables, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the limit is exceeded, you can add synchronization objects in batches by editing the synchronization objects. For details, see Editing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-32 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-33 Task startup settings

Table 4-30 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.4 From MySQL to Oracle

Supported Source and Destination Databases

Table 4-31 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> RDS for MySQL (5.5, 5.6, 5.7, and 8.0) 	<ul style="list-style-type: none"> On-premises databases ECS databases

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-level comparison
To obtain accurate comparison results, **compare data** at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-32 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. • The destination database user must have the following permissions: ALTER ANY INDEX, ALTER ANY TABLE, ALTER SESSION, ANALYZE ANY, COMMENT ANY TABLE, CREATE ANY INDEX, CREATE ANY TABLE, CREATE SESSION, DELETE ANY TABLE, DROP ANY TABLE, INSERT ANY TABLE, SELECT ANY TABLE, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, UPDATE ANY TABLE, and RESOURCE roles.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Full synchronization supports the synchronization of data, table structures, and indexes. ● Incremental synchronization supports only data synchronization. ● Geography data types such as geometry, geometrycollection, linestring, multilinestring, multipoint, point and polygon are not supported. ● Views, constraints, functions, stored procedures, triggers, and events cannot be synchronized. ● The system database and event status cannot be synchronized. ● The destination Oracle database does not support empty strings, so the object to be synchronized cannot contain empty strings. ● The maximum number of columns supported by the source MySQL database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● The source database names cannot contain non-ASCII characters, or the following characters: <code>:'<>^\"</code> ● The table name in the source database cannot contain non-ASCII characters or the following characters: <code>:'<>^\"</code> ● The column names in the source database tables cannot end with a backslash (<code>\</code>). ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails.

Type	Restrictions
	<ul style="list-style-type: none">• The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.• The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail.• If the source MySQL database version is 8.0, do not set lower_case_table_names to 0.• The source database cannot be a read replica.
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly.• The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● The NOT NULL constraint of MySQL supports empty strings, while the NOT NULL constraint of Oracle does not. During a synchronization, if an empty string exists in a NOT NULL constraint field, delete the NOT NULL constraint from the destination Oracle database. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the default value of the time field in the source database is all 0s, it will be converted to 1970-01-01 00:00:00. ● If the precision of source database decimal data type exceeds 38, the data will be truncated because the maximum precision of the destination database number data type is 38. ● The varbinary, binary, and timestamp columns in the source database cannot contain primary key or unique constraints. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● The source database does not support fields whose column type is binary and the length is 0 (that is, binary(0)). This is because the binary type of MySQL is mapped to the raw type when being synchronized to Oracle using DRS. However, in Oracle, the length of the raw type cannot be set to 0. ● All table field names are converted to uppercase letters. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If

Type	Restrictions
	<p data-bbox="667 297 1422 427">newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> <li data-bbox="592 443 1390 510">● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. <li data-bbox="592 521 1426 618">● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. <li data-bbox="592 629 1426 725">● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. <li data-bbox="592 736 1417 808">● During the synchronization, the source database cannot write data using the statement-based binlog format. <li data-bbox="592 819 1426 887">● During the synchronization, do not clear binlogs on the source database. <li data-bbox="592 898 1378 994">● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. <li data-bbox="592 1005 1406 1077">● During the synchronization, do not create a database named ib_logfile in the source. <li data-bbox="592 1088 1385 1184">● During an incremental synchronization, do not perform the point-in-time recovery (PITR) operation on the source database. <li data-bbox="592 1196 1410 1267">● During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. <li data-bbox="592 1279 1374 1312">● Incremental synchronization filters out all DDL operations. <li data-bbox="592 1323 1406 1453">● During incremental synchronization, resumable upload is supported, but data may be repeatedly inserted into non-transactional tables that do not have primary keys when the server system breaks down. <li data-bbox="592 1464 1362 1536">● If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. <li data-bbox="592 1547 1426 1644">● Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-34 Synchronization task information

Warning: Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page. The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-5678

Description: [Text area]

Table 4-33 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-35 Synchronization instance details

Synchronization Instance Details

Data Flow: To the cloud | **Out of the cloud** | Self-built to self-built

Source DB Engine: MySQL | EDM | GaussDB Distributed | GaussDB Primary/Standby | MariaDB | COS | PostgreSQL | GaussDB for MySQL

Destination DB Engine: MySQL | **Click** | CSSECS | Kafka | MemDB

Network Type: Public network

Source DB Instance: No DB instance available | View DB Instance | View Unavailable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnets | View Occupied IP Address

Synchronization Mode: Full synchronization

Specify EP: [Dropdown menu] | Create an EP

Table 4-34 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .

Parameter	Description
Destination DB Engine	Select Oracle .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-36 Task type




Table 4-35 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-37 Enterprise Project and Tags




Table 4-36 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-38 Source database information

Source Database

DB Instance Name

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-37 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-39 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ?

Port

Database Service Name Service Name

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-38 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>The IP address or domain name of the destination database.</p> <p>NOTE</p> <p>For a RAC cluster, use a scan IP address to improve access performance.</p>
Port	The port of the destination database. Range: 1 - 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-40 Synchronization mode

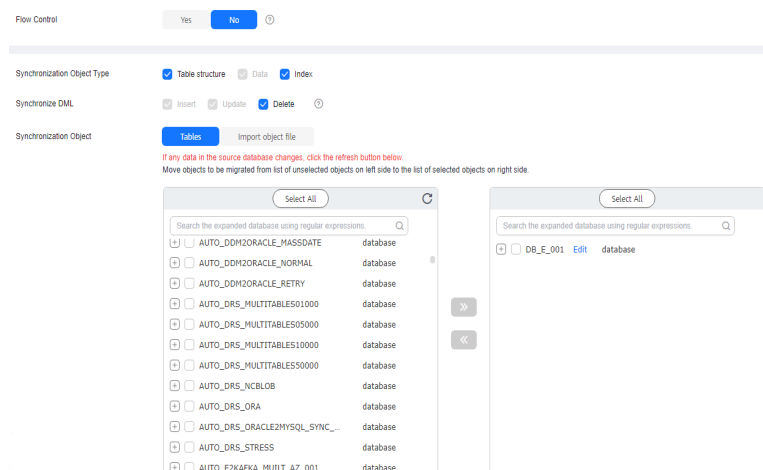
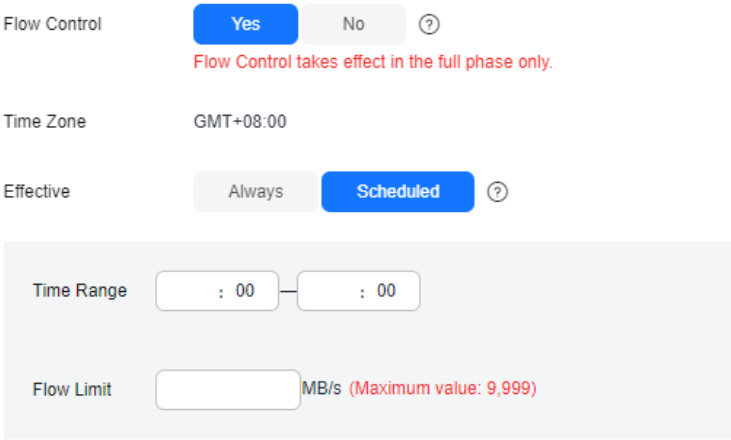



Table 4-39 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-41 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <p>For details about how to import an object file, see Importing Synchronization Objects.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-42 Task startup settings

Table 4-40 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.5 From MySQL to MariaDB

Supported Source and Destination Databases

Table 4-41 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> RDS for MySQL 5.6, 5.7, and 8.0 	<ul style="list-style-type: none"> On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 <p>NOTE</p> <ul style="list-style-type: none"> If the source database version is MySQL 5.6, you are advised to select MariaDB 10.0, 10.1, or later as the destination database. If the source database version is MySQL 5.7, you are advised to select MariaDB 10.2, 10.3, 10.4, or later as the destination database. If the source database version is MySQL 8.0, you are advised to select MariaDB 10.5 as the destination database.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-42](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 4-42 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user	The user must have the following minimum permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 4-43](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-43 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. • Events and triggers cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 4-44](#).

Table 4-44 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: .></\'" • The column names in the source database tables cannot end with a backslash (\). • The table names of the source database cannot contain non-ASCII characters, or the following special characters: ></\ • The source database names cannot contain non-ASCII characters, or the following characters: . > < / \ ' `" • The table names of the source database cannot contain non-ASCII characters, or the following special characters: > < / \ • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • If the source MySQL database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● A real-time synchronization task may fail due to the change of the username and password of the source or destination database. If it happens, rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization. ● To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization. ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> ● When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● Modifying MyISAM tables may cause data inconsistency. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● If the session variable character_set_client is set to binary, some data may include garbled characters. ● Some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out.

Type	Restrictions
	<ul style="list-style-type: none">- If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out.- You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.• During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization?• You can add additional synchronization objects. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the enforce_storage_engine parameter of the destination MariaDB database is set to InnoDB, DRS cannot synchronize the table structure and data whose storage engine is MyISAM to the destination MariaDB database. To synchronize table data whose storage engine is MyISAM, create a table structure on the destination database. (The storage engine can only be set to InnoDB due to the value restriction of the enforce_storage_engine parameter.) ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions? ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The source database does not support point-in-time recovery (PITR).

Type	Restrictions
	<ul style="list-style-type: none"> ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The partitioned table does not support column mapping. ● After a task is created, the destination database cannot be set to read-only. ● During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. ● If the source database version is MySQL 8.0, some collation character sets (such as utf8mb4_0900_as_ci, utf8mb4_0900_as_cs, utf8mb4_0900_bin and utf8mb4_cs_0900_ai_ci that support Unicode 9.0) are not supported by the destination database. ● Due to engine and version differences, the following MySQL functions are not supported in MariaDB: MBRCOVEREDBY, ST_BUFFER_STRATEGY, ST_GeoHash, ST_IsValid, ST_LatFromGeoHash, ST_LongFromGeoHash, ST_PointFromGeoHash, ST_SIMPLIFY, ST_VALIDATE, (8.0)JSON_ARRAYAGG, JSON_OBJECTAGG, JSON_PRETTY, JSON_STORAGE_FREE, JSON_STORAGE_SIZE and JSON_TABLE. ● The destination MariaDB database does not support VALIDATION during partition exchange. Do not use the verification syntax when executing partition exchange SQL statements in the source MySQL database. For example: ALTER TABLE t1 EXCHANGE PARTITION p0 WITH TABLE t2 WITH VALIDATION; ● Do not write characters such as \n, \t, and \r to the JSON data of the source MySQL database during incremental synchronization. Otherwise, there may be data inconsistencies. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

This section describes how to synchronize data out of the cloud from an RDS for MySQL database to a MariaDB database. To configure other storage engines, you can refer to the following procedures.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-43 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ 0/256

Table 4-45 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 4-44 Synchronization instance information

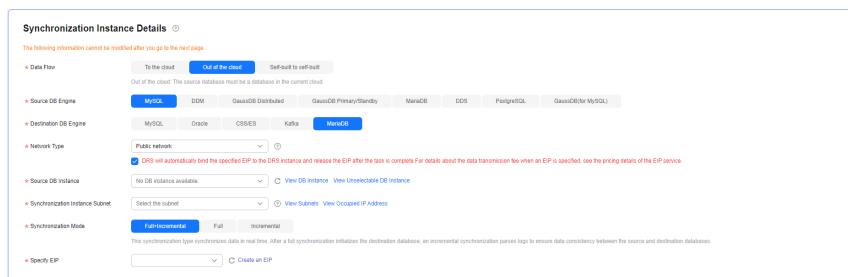


Table 4-46 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MariaDB .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the source DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-45 AZ





Table 4-47 Task AZ

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-46 AZ</p> 

- Enterprise Project and Tags

Figure 4-47 Enterprise Project and Tags




Table 4-48 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.


- Source database information

Figure 4-48 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-49 Source database information

Parameter	Description
DB Instance Name	The RDS for MySQL instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.

Parameter	Description
Database Password	The password for the database username.
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

- Destination database information

Figure 4-49 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 4-50 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - Only .cer and .pem certificates are supported. - If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 4-50 Synchronization Mode

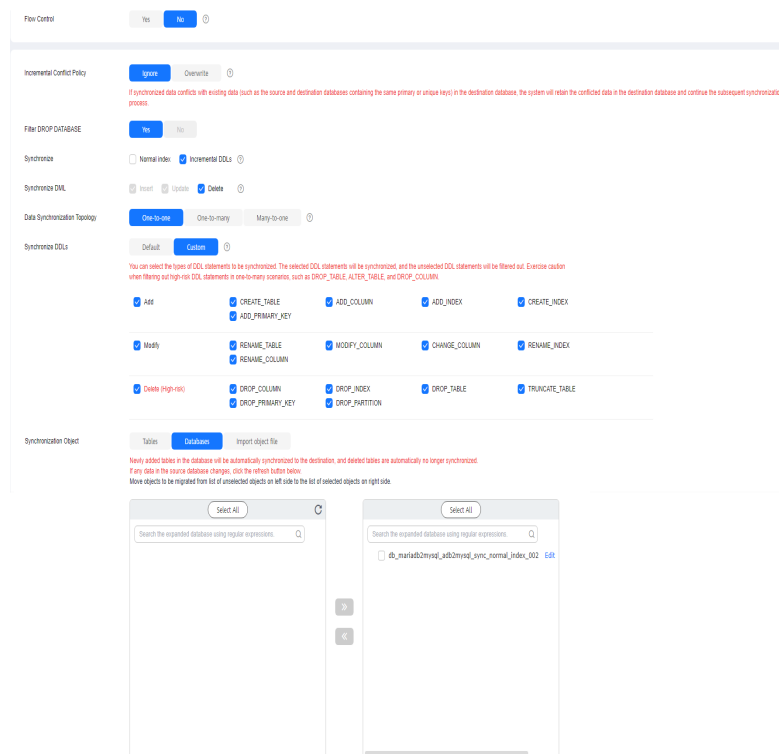
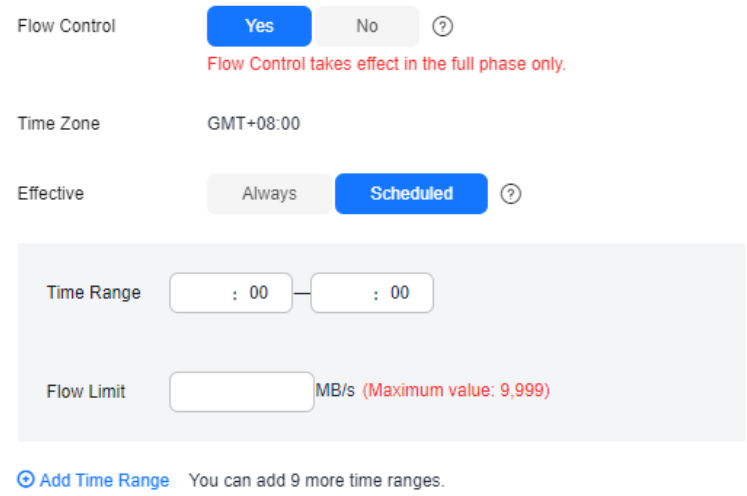



Table 4-51 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-51 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>


Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-52 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

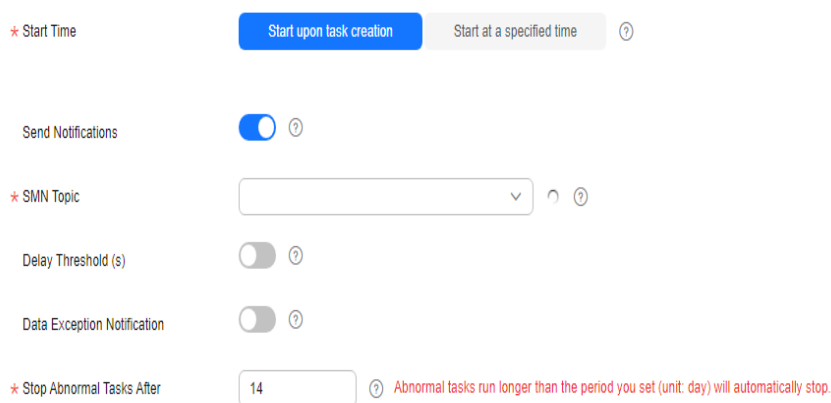
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-53 Task startup settings



* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ?

Delay Threshold (s) ?

Data Exception Notification ?


* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 4-52 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.6 From DDM to MySQL

Supported Source and Destination Databases

Table 4-53 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> DDM instances 	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-54](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 4-54 Database account permission

Type	Full, Full+Incremental, and Incremental
Source database user	<ul style="list-style-type: none"> The user of the source DDM database must have at least one permission, for example, SELECT. The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions.
Destination database user	<p>The user must have the SELECT, CREATE, ALTER, DROP, DELETE, INSERT and UPDATE permissions.</p> <p>The root account of the RDS for MySQL DB instance has the preceding permissions by default.</p>

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 4-55 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-55 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Full synchronization supports the synchronization of data, table structures, and indexes.• The source database cannot contain tables whose sharding keys are of the timestamp type.• The sharding key of the source table must be added to the primary key and unique key of the destination table, which means that the primary key and unique key columns of the destination table must contain the sharded columns of the source table to avoid data conflict and inconsistency.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see **Pre-check Items**. In addition to the pre-check items, you need to pay attention to the items listed in **Table 4-56**.

Table 4-56 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: .'<>\ • Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. <p>Full synchronization</p> <p>During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket.</p> <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not perform data restoration on the source database. Some DDL operations are supported. <ul style="list-style-type: none"> DROP_DATABASE, DROP_TABLE, TRUNCATE_TABLE, CREATE_VIEW and DROP_VIEW are not supported. Online DDL is not supported. Tables can be created. For example: <pre>create table `ddl_test` (id int, c1 varchar(25), primary key(id)); create table `ddl_test_gho` like `ddl_test`;</pre> Tables can be renamed. Both the source and destination tables must be selected. For example: <pre>rename table `ddl_test` to `ddl_test_new`;</pre> Columns in a table can be added and modified, but cannot be deleted. For example: <pre>alter table `ddl_test` add column `c2` varchar(25); alter table `ddl_test` modify column `c1` varchar(50); alter table `ddl_test` alter c1 set default '***';</pre> Table indexes can be modified. For example: <pre>alter table `ddl_test` drop primary key; alter table `ddl_test` add primary key(id); alter table `ddl_test` add index `ddl_test_uk`(id); alter table `ddl_test` drop index `ddl_test_uk`;</pre> In table-level synchronization, you can add columns, modify columns, and add primary keys and normal indexes.

Type	Restrictions
	<ul style="list-style-type: none"> - During database-level synchronization, you can create tables, rename tables, add columns, modify columns, and add primary keys and normal indexes. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. • Perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. • The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● Resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down. ● If the source database contains a duplicate primary key or unique key, the data synchronized to the destination database will be less than that in the source database. Therefore, you must check and correct the data before starting the synchronization task. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If a table in the source database contains system columns <code>_ddm_lock</code> and <code>_slot</code>, data in these columns will be filtered out during full synchronization and incremental synchronization, resulting in data loss. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, the destination database cannot be set to read-only. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-54 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.

The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region ▼

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project ▼

* Task Name 🔍

DRS-5678

Description 🔍

0/256

Table 4-57 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 4-55 Synchronization instance details

Synchronization Instance Details 🔍

⚠ The following information cannot be modified after you go to the next page.

Data Flow 🔍

To the cloud: Out of the cloud Self-built to self-built

Out of the cloud: This scenario requires replication to a database in the current cloud.

Source DB Engine 🔍

MySQL ODM GaussDB Distributed GaussDB Primary/Standby MariaDB DCS PostgreSQL GaussDB for MySQL

Destination DB Engine 🔍

MySQL Oracle Kafka

Network Type 🔍

Public network

DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

Source DB Instance 🔍

No DB instance available 🔍 View DB Instance 🔍 View Unavailable DB Instance

Synchronization Instance Subnet 🔍

Select the subnet 🔍 View Subnets

Synchronization Mode 🔍

Full synchronization Full Incremental

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Source DB Instance Quantity 🔍

1

The number of source DB instances must be the same as the number of instances associated with DRS.

Specify EP 🔍

Create an EP

Table 4-58 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select MySQL .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The DDM instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Quantity	Specifies the number of DB instances bound to the source DDM database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-56 Task type




Table 4-59 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-57 Enterprise Project and Tags




Table 4-60 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-58 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-61 Source database settings

Parameter	Description
DB Instance Name	The DDM instance you selected when you create a synchronization task. The instance name cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 4-59 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 4-62 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-60 Synchronization Mode

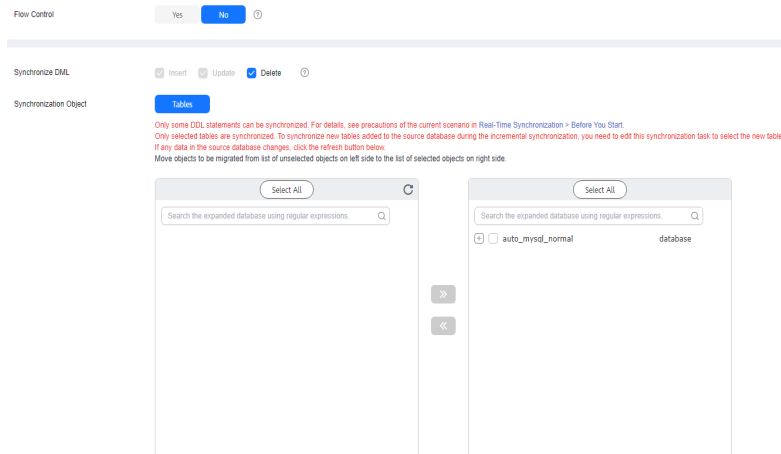
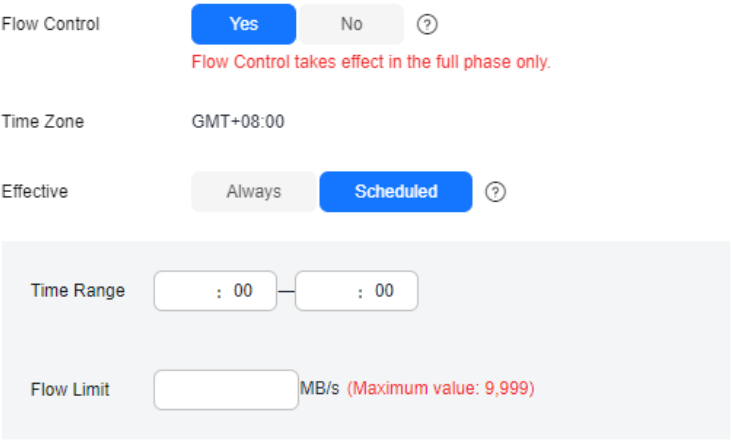



Table 4-63 Synchronization mode and object

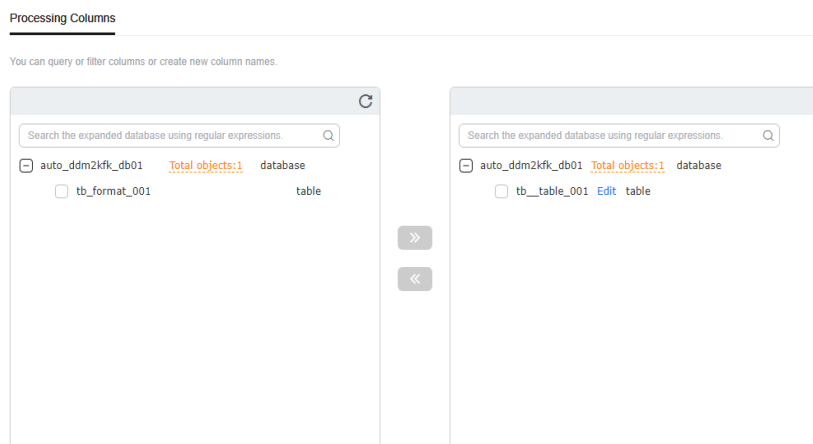
Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-61 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 4-62 Data processing



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-63 Task startup settings





Table 4-64 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.7 From DDM to Oracle

Supported Source and Destination Databases

Table 4-65 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DDM instances 	<ul style="list-style-type: none"> • On-premises databases • ECS databases

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-66 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The destination database user must have the following permissions: ALTER ANY INDEX, ALTER ANY TABLE, ALTER SESSION, ANALYZE ANY, COMMENT ANY TABLE, CREATE ANY INDEX, CREATE ANY TABLE, CREATE SESSION, DELETE ANY TABLE, DROP ANY TABLE, INSERT ANY TABLE, SELECT ANY TABLE, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, UPDATE ANY TABLE, and RESOURCE roles.
Synchronization object	<ul style="list-style-type: none"> • The source database data can be synchronized. • The source database table structure supports only full synchronization. • Database objects other than table structures, indexes, and constraints cannot be synchronized. • The destination Oracle database does not support empty strings, so the object to be synchronized cannot contain empty strings. • The maximum number of columns supported by the source DDM database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During the incremental synchronization, the binlog of the source sharded database must be enabled and use the row-based format. ● If the storage space is sufficient, store the binlogs for as long as possible. The recommended retention period is three days. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: <code>!<>/\</code> ● Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout. ● Enable GTID on the source sharded database. ● The source physical sharded database does not support the enumeration and set types. ● The source database table name and field name cannot exceed 30 characters. ● Tables without primary keys cannot be synchronized. ● The source databases to be synchronized must have the RESOURCE permission. ● The default value in the timestamp column of the source database must be within the range allowed by the destination database. Otherwise, the synchronization fails.
Destination database	<ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. ● The time zone of the destination database must be the same as that of the source database. ● The destination database (account) must have the RESOURCE permission.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The NOT NULL constraint of MySQL supports empty strings, while the NOT NULL constraint of Oracle does not. During a synchronization, if an empty string exists in a NOT NULL constraint field, delete the NOT NULL constraint from the destination Oracle database. ● If the source database contains a duplicate primary key, data synchronized to the destination database will be less than that in the source database. You must check and correct the data before starting the synchronization task. ● The varbinary, binary, and timestamp columns in the source table cannot contain primary key or unique constraints. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● If a table in the source database contains system columns _ddm_lock and _slot, data in these columns will be filtered out during full synchronization and incremental synchronization, resulting in data loss. ● If the sharding key of a sharded database or table is a function rather than a primary key, the primary key and sharding key will be automatically used as the composite primary key when the table is created in Oracle. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. ● DDL operations are not supported during synchronization. ● During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an

Type	Restrictions
	<p>unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table.</p> <ul style="list-style-type: none"> • During an incremental synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During an incremental synchronization, if you need to modify the structure of the source table to be synchronized, you must modify the corresponding destination table structure. • During an incremental synchronization, do not perform the restoration operation on the source database.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-64 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-67 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

Parameter	Description
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 4-65 Synchronization instance details

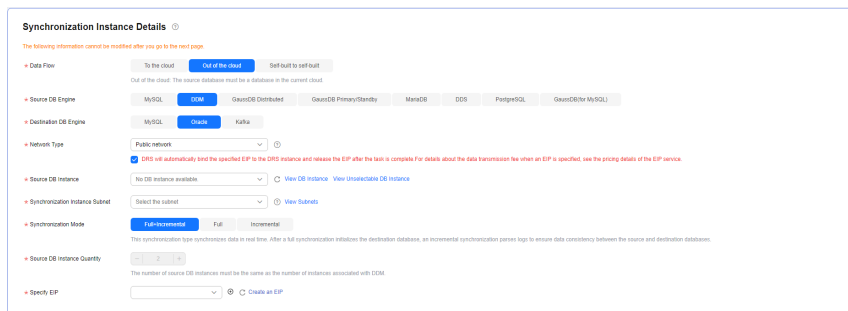


Table 4-68 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select Oracle .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The DDM instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Quantity	<p>Specifies the number of DB instances bound to the source DDM database.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-66 Task type




Table 4-69 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-67 Enterprise Project and Tags




Table 4-70 Enterprise Project and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .

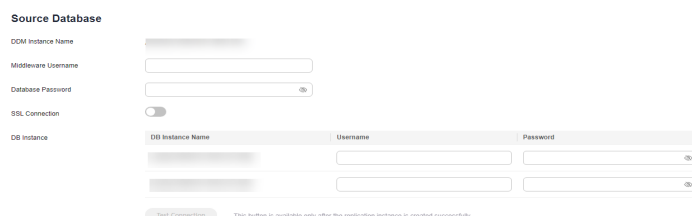
Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-68 Source database information



The screenshot shows the 'Source Database' configuration interface. It includes input fields for 'DDM Instance Name', 'Middleware Username', and 'Database Password'. There is a toggle for 'SSL Connection'. Below this is a table for 'DB Instance' configuration with columns for 'DB Instance Name', 'Username', and 'Password'. A 'Test Connection' button is located at the bottom, with a note: 'This button is available only after the replication instance is created successfully.'

Table 4-71 Source database settings

Parameter	Description
DDM Instance Name	The DDM instance you selected when you create a synchronization task. The instance name cannot be changed.
Middleware Username	The username for accessing the source database.
Database Password	The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.
DB Instance	The sharded database details.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 4-69 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

Database Username

Database Password ⓘ

SSL Connection

Table 4-72 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>The IP address or domain name of the destination database.</p> <p>NOTE</p> <p>For a RAC cluster, use a scan IP address to improve access performance.</p>
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-70 Synchronization mode

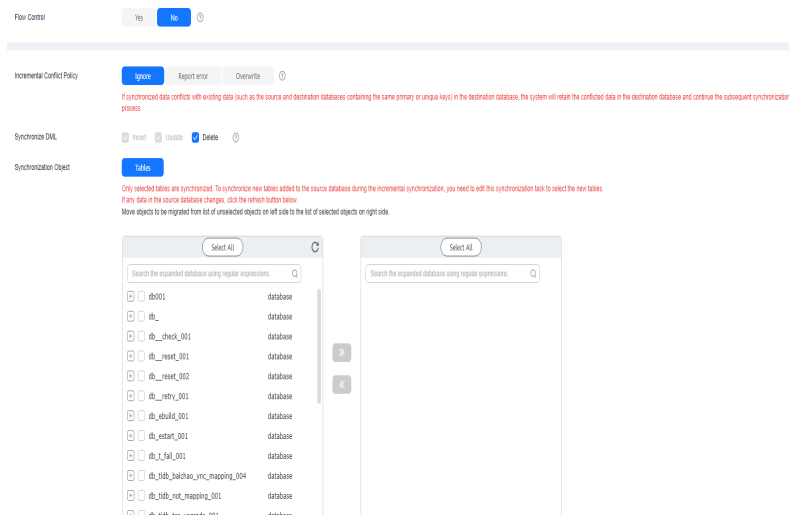
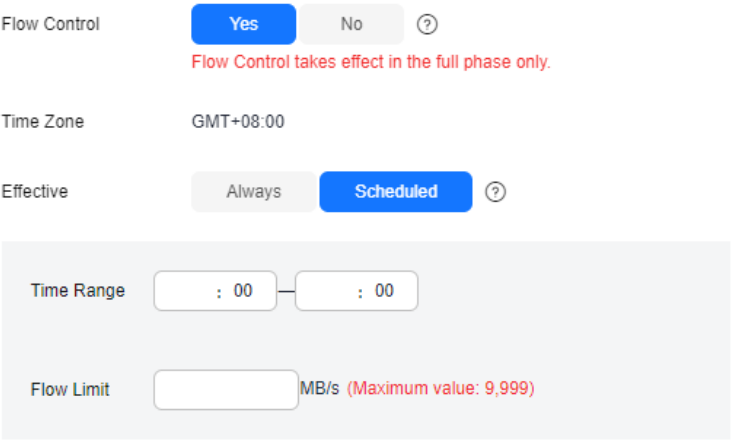



Table 4-73 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-71 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Data Processing** page, select the table object to be processed, enter the column name, type, and operation type to be added, confirm the information, and click **Next**. You can configure related rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

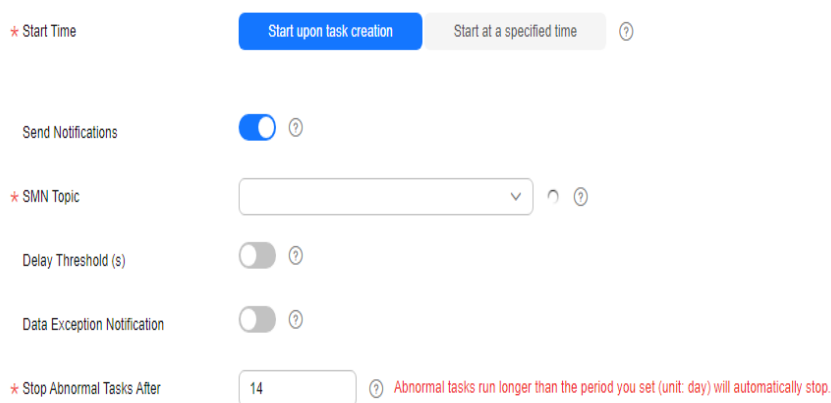
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-72 Task startup settings



* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ?

Delay Threshold (s) ?

Data Exception Notification ?


* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 4-74 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.8 From DDM to Kafka

Supported Source and Destination Databases

Table 4-75 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> DDM instances 	<ul style="list-style-type: none"> Kafka

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-76 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT.
Synchronization object	<ul style="list-style-type: none"> The table data can be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During the incremental synchronization, the binlog of the source physical sharded database must be enabled and in row format. Enable GTID and set the value of binlog_row_image to FULL. ● If the storage space is sufficient, store the binlogs for as long as possible. The recommended retention period is three days. ● The source database server_id must be set. The value of server_id ranges from 1 to 4294967296. ● The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: <code>.'<>/\</code> ● Enable skip-name-resolve for the source sharded database to reduce the possibility of connection timeout. ● Enable GTID on the source sharded database. ● The source sharded database does not support real-time synchronization of the enumerated and set types. ● The source database table name and field name cannot exceed 30 characters.
Destination database	<ul style="list-style-type: none"> ● The destination database is a Kafka database. ● You are advised to set auto.create.topics.enable of Kafka to false.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> • If the data types are incompatible, the synchronization may fail. • If the source database contains a duplicate primary key, data synchronized to the destination database will be less than that in the source database. Therefore, you must check and correct the data before starting the synchronization task. • If a table in the source database contains system columns <code>_ddm_lock</code> and <code>_slot</code>, data in these columns will be filtered out during data synchronization, resulting in data loss. • After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. • DDL operations are supported during incremental synchronization. • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-73 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 4-77 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&""

- Synchronization instance details

Figure 4-74 Synchronization instance details

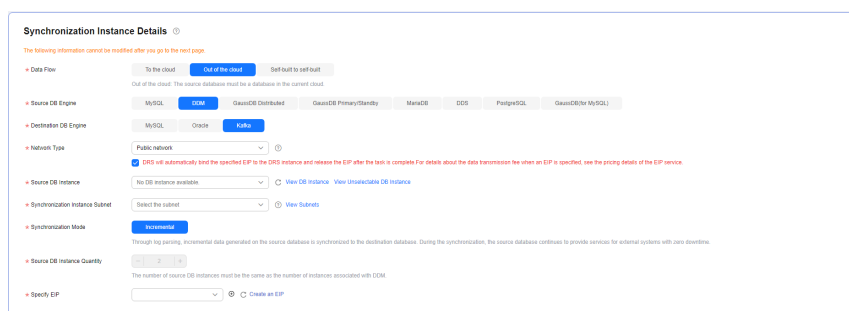


Table 4-78 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select Kafka .

Parameter	Description
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	An available DDM instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>
Source DB Instance Quantity	Specifies the number of DB instances bound to the source DDM database. Set this parameter based on the site requirements.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-75 Task type




Table 4-79 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-76 Enterprise Project and Tags




Table 4-80 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-77 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-81 Source database settings

Parameter	Description
DB Instance Name	The DDM instance you selected when you create the synchronization task. The instance name cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 4-78 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

✔ Test successful

Table 4-82 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-79 Synchronization mode

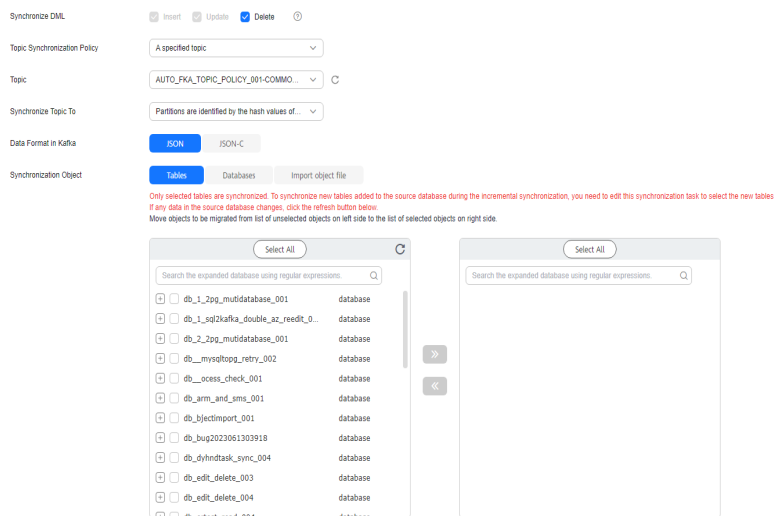



Table 4-83 Synchronization Object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .

Parameter	Description
Topic Name Format	<p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name.</p> <p>For example, if this parameter is set to \$database\$-tablename\$ and the database name is db1, and the table name is tab1, the topic name is db1-tab1. If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic.
Data Format in Kafka	<p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

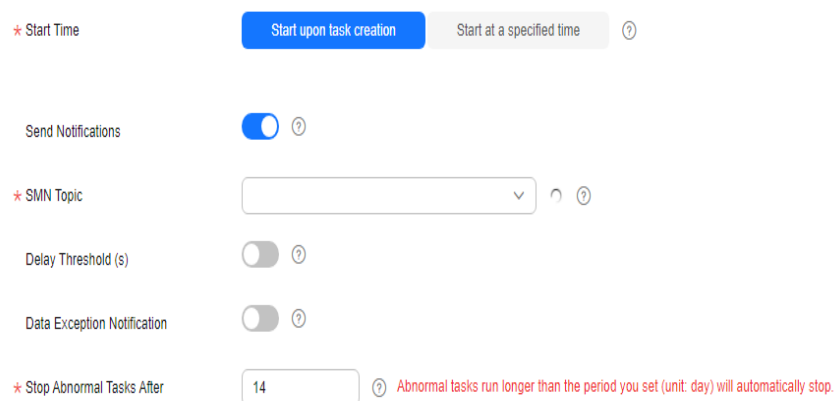
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-80 Task startup settings




The screenshot shows a configuration form for task startup settings. It includes the following fields and controls:

- * Start Time:** Two radio buttons: "Start upon task creation" (selected) and "Start at a specified time".
- Send Notifications:** A toggle switch that is turned on.
- * SMN Topic:** A dropdown menu with a search icon and a help icon.
- Delay Threshold (s):** A toggle switch that is turned off.
- Data Exception Notification:** A toggle switch that is turned off.
- * Stop Abnormal Tasks After:** A text input field containing the value "14". A red note below it states: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-84 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.9 From DDS to MongoDB

Supported Source and Destination Databases

Table 4-85 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, 4.4 and 5.0) <p>NOTE DDS 5.0 supports replica sets only.</p>	<ul style="list-style-type: none"> • On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) • ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) • Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) <p>NOTE The destination database version must be the same as or later than the source database version.</p>

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-86 Precautions

Type	Restrictions
Database permissions	<p>Source database (minimum permissions):</p> <ul style="list-style-type: none"> • Permission requirements for incremental synchronization: <ul style="list-style-type: none"> – Replica set: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. – Single node: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. – Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. <p>Minimum permission requirements: The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the read permission for the config database.</p>

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. ● Cluster: Currently, only collections (including validator and capped and non-capped collections), shard keys, indexes, and views can be synchronized. ● Single node: Only collections (including validator and capped collections), indexes, and views can be synchronized. ● Collections that contain the <code>_id</code> field without indexes are not supported. ● The first parameter of <code>BinData()</code> cannot be 2. ● If ranged sharding is used, <code>maxKey</code> cannot be used as the primary key. ● If the source database is a cluster database, <code>DBPointer</code> and <code>DBRef</code> are not supported. ● Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization. ● Time series collections are not supported.
Source database	<ul style="list-style-type: none"> ● During the incremental synchronization, the Olog of the source database must be enabled. ● If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is three days. ● The source database name cannot contain <code>\.</code> or spaces. The collection name and view name cannot start with system. or contain the dollar sign (<code>\$</code>). ● If the source database is a cluster, the balancer must be disabled and orphan documents must be deleted. For details, see How Do I Disable the Balancer? and How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?
Destination database	<ul style="list-style-type: none"> ● Data cannot be synchronized from a newer version database to an older version database. ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the source is not a cluster instance, the following operations and commands are supported during incremental synchronization: <ul style="list-style-type: none"> - Creating and deleting databases - Adding, deleting, and updating documents - Creating and deleting collections - Creating and deleting indexes - Creating and deleting views - The <code>convertToCapped</code>, <code>collMod</code>, and <code>renameCollection</code> commands are supported. ● If the source is a cluster instance: <ul style="list-style-type: none"> - The cluster version cannot be earlier than 4.0. - Do not delete synchronization objects in the incremental synchronization phase. Otherwise, the synchronization task will fail. - The synchronization process consumes certain CPU and memory resources of the source database. Evaluate the source database resources in advance. - If the load on the source database is heavy, the processing speed of change streams cannot keep up with the oplog generation speed. As a result, DRS synchronization delay occurs. - Only the following DDLs are supported: drop database, drop collection and rename - In the incremental synchronization phase, the synchronization speed can reach up to 10,000 rows in a single table per second. ● During startup, the specified start point must be within the Oplog range. ● The destination database user must have the write permission. If the destination is a cluster instance, the database user must have the read permission for the config data. ● If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental synchronization, data consistency cannot be ensured when source and destination databases are in different time zone. ● The value of <code>block_compressor</code> is determined by <code>stats().wiredTiger.creationString.block_compressor</code> of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be migrated. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on <code>net.compression.compressors</code> of the destination

Type	Restrictions
	<p>database. If the storage engine of the destination database is not WiredTiger, DRS does not synchronize compression parameters.</p> <ul style="list-style-type: none"> ● If the destination is a replica set instance out of the cloud, enter information about all primary and secondary nodes to reduce the impact of a primary/secondary switchover on the synchronization task. If you enter information about primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● If the destination is a cluster instance out of the cloud, enter information about multiple mongos nodes to reduce the impact of the single-node failure on the synchronization task. In addition, multiple mongos nodes support load balancing. In addition, ensure that all mongos nodes belong to the same cluster instance. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during the entire synchronization process. ● During the synchronization, data rollback caused by a primary/standby switchover of the source database is not supported. ● The system database cannot be synchronized. The username and role must be manually created in the destination database. ● Documents larger than 16 MB in the source database cannot be inserted or updated during incremental synchronization. ● In the incremental synchronization phase, concurrent replay is performed at the collection level to maintain the synchronization performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. <p>In either of the preceding scenarios, the task delay may increase.</p> ● To prevent loopback, the to-the-cloud migration task and out-of-cloud synchronization task cannot be started at the same time. ● During row comparison, if an orphan document exists in a cluster instance or chunks are being synchronized, the number of returned rows is incorrect and the comparison results are inconsistent. For details, see MongoDB official documentation.

Procedure

This section uses real-time incremental synchronization from DDS to MongoDB as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-81 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-87 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-82 Synchronization instance details

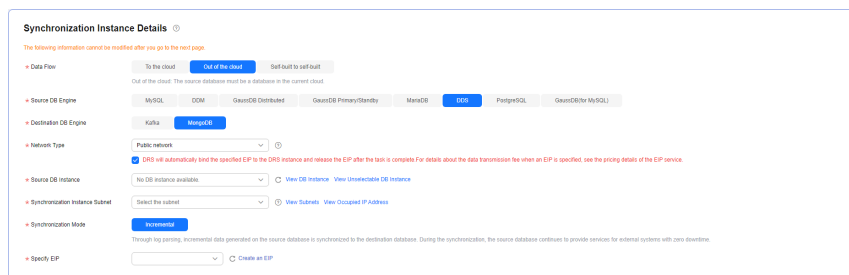


Table 4-88 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDS .
Destination DB Engine	Select MongoDB .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	An available DDS instance
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-83 Task type




Table 4-89 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p>

- Enterprise Project and Tags

Figure 4-84 Enterprise Project and Tags




Table 4-90 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-85 Source database information

Source Database

DB Instance Name

Authentication Database

Database Username

Database Password

SSL Connection

Table 4-91 Source database settings

Parameter	Description
DB Instance Name	The DDS instance you selected when creating the task. This parameter cannot be changed.
Authentication Database	The name of the authentication database. For example: The default authentication database of Huawei Cloud DDS instance is admin .
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 4-86 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Authentication Database

Database Username

Database Password

SSL Connection

Test Connection This button is available only after the replication instance is created successfully.

Table 4-92 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <p>You can enter up to 3 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080. Ensure that the entered IP addresses or domain names belong to the same instance.</p> <p>NOTE If multiple IP addresses or domain names are entered, the test connection is successful as long as one IP address or domain name is accessible. Therefore, you must ensure that the IP address or domain name is correct.</p>
Authentication Database	The name of the destination database.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Task** page, select the start point and synchronization objects, and click **Next**.

Figure 4-87 Synchronization mode

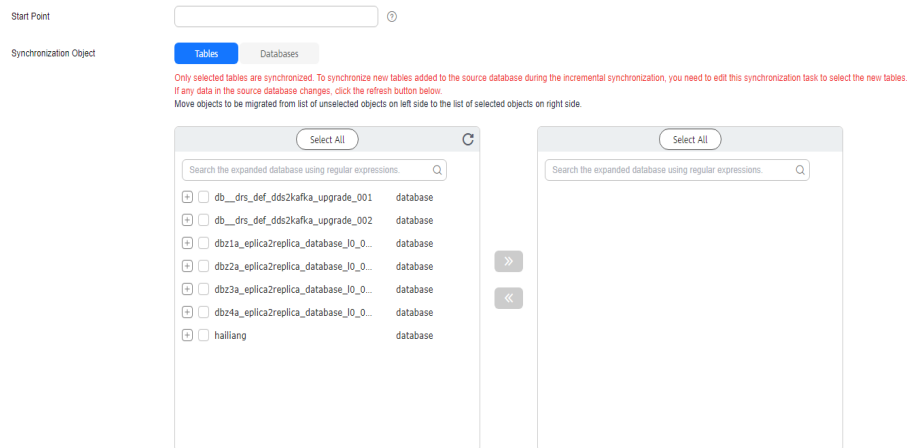



Table 4-93 Synchronization Object

Parameter	Description
Start Point	<p>The start time of incremental synchronization. The value is in the format of timestamp:incr. timestamp is the Unix timestamp (unit: second), and incr is the command execution sequence in a second. A synchronization task obtains incremental logs of the source database from the start position (including the current start position).</p> <ul style="list-style-type: none"> Run <code>db.getSiblingDB("local").oplog.rs.find()</code> to query the source database oplog. The format of the <code>ts</code> field is timestamp:incr. A cluster cannot query the oplog from mongos. timestamp can be converted into the Unix timestamp format based on the start time. The value of incr is 1. <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

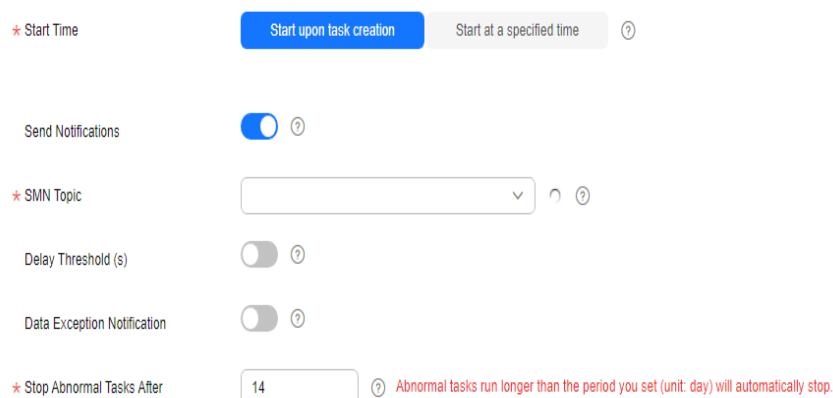
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-88 Task startup settings




The screenshot shows a configuration form for task startup settings. It includes the following fields and controls:

- * Start Time:** Two radio buttons are present. The first, labeled "Start upon task creation", is selected and highlighted in blue. The second, labeled "Start at a specified time", is unselected. A help icon (?) is to the right.
- Send Notifications:** A toggle switch is turned on (blue). A help icon (?) is to the right.
- * SMN Topic:** A dropdown menu is currently empty. A refresh icon and a help icon (?) are to the right.
- Delay Threshold (s):** A toggle switch is turned off (grey). A help icon (?) is to the right.
- Data Exception Notification:** A toggle switch is turned off (grey). A help icon (?) is to the right.
- * Stop Abnormal Tasks After:** A text input field contains the value "14". A help icon (?) is to the right, followed by a red warning message: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-94 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.10 From DDS to Kafka

Supported Source and Destination Databases

Table 4-95 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DDS DB instances (versions 4.0, 4.2, 4.4 and 5.0) <p>NOTE DDS 5.0 supports replica sets only.</p>	<ul style="list-style-type: none"> • Kafka

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-96 Precautions

Type	Restrictions
Database permissions	Source database (minimum permissions): <ul style="list-style-type: none"> • Permission requirements for incremental synchronization: The source database user must have the readAnyDatabase permission for the admin database.
Source database	<ul style="list-style-type: none"> • During the incremental synchronization, the Oplog of the source database must be enabled. • If the storage space is sufficient, store the source database Oplog for as long as possible. The recommended retention period is three days. • The source database name cannot contain \\. "\$ or spaces. The collection name cannot start with system. or contain the dollar sign (\$). • If the source database is a cluster, the balancer must be disabled and orphan documents must be deleted. For details, see How Do I Disable the Balancer? and How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?
Destination database	<ul style="list-style-type: none"> • The destination database is a Kafka database. • You are advised to set auto.create.topics.enable of Kafka to false.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> • The synchronization process consumes certain CPU and memory resources of the source database. Evaluate the source database resources in advance. • If the source database is under a heavy load, the processing speed of change streams cannot keep up with the oplog generation speed. As a result, DRS synchronization delay occurs. • The following DMLs are supported: insert, update, replace, and delete. • Only the following DDLs are supported: drop database, drop collection and rename • In the incremental synchronization phase, the synchronization speed can reach up to 10,000 rows in a single table per second. • During startup, the specified start point must be within the Oplog range. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during the entire synchronization process. • Documents larger than 16 MB in the source database cannot be inserted or updated during incremental synchronization. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail.

Procedure

This section uses real-time incremental synchronization from DDS to Kafka as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-89 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 4-97 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 4-90 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud **Out of the cloud** Self-host to self-host
Out of the cloud: The source database resides in a database in the current cloud.

Source DB Engine: MySQL, DDS, GreenDB Distributed, GreenDB Primary/Standby, MariaDB, **DDS**, PostgreSQL, GreenDB for MySQL

Destination DB Engine: **Kafka**, MongoDB

Network Type: ⓘ

Source DB Instance: ⓘ View DB Instance View Unavailable DB Instance

Synchronization Instance Subnet: ⓘ View Subnets View Occupied IP Address

Synchronization Mode: **Incremental**
Through log parsing, incremental data generated on the source database is synchronized to the destination database. During the synchronization, the source database continues to provide services for external systems with zero downtime.

Specify EP: ⓘ Create an EP

Table 4-98 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDS .
Destination DB Engine	Select Kafka .

Parameter	Description
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	An available DDS instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-91 Task type




Table 4-99 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-92 Enterprise Project and Tags




Table 4-100 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-93 Source database information

Source Database

DB Instance Name

Authentication Database

Database Username

Database Password

SSL Connection

Table 4-101 Source database settings

Parameter	Description
DB Instance Name	The DDS instance you selected when creating the task. This parameter cannot be changed.
Authentication Database	The name of the authentication database. For example: The default authentication database of Huawei Cloud DDS instance is admin .
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 4-94 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

 Test successful

Table 4-102 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

 NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Task** page, select the start point and synchronization objects, and click **Next**.

Figure 4-95 Synchronization Mode

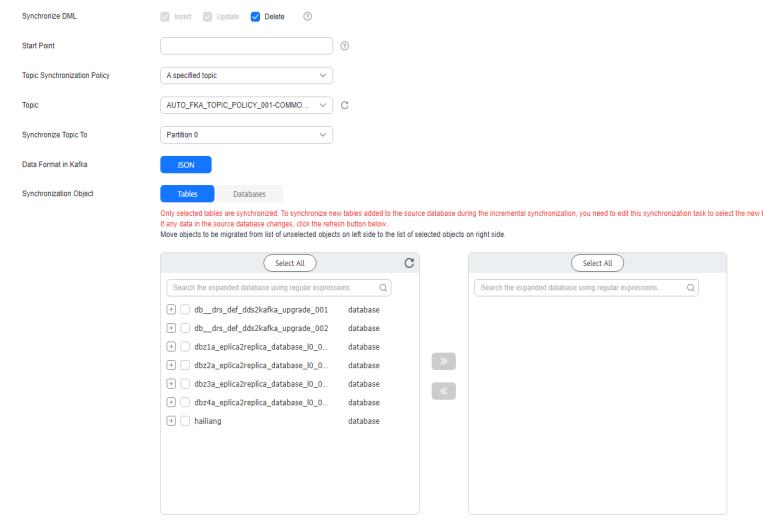



Table 4-103 Synchronization Object

Parameter	Description
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>The start time of incremental synchronization. The value is in the format of timestamp:incre. timestamp is the Unix timestamp (unit: second), and incre is the command execution sequence in a second. A synchronization task obtains incremental logs of the source database from the start position (including the current start position).</p> <ul style="list-style-type: none"> Run db.getSiblingDB("local").oplog.rs.find() to query the source database oplog. The format of the ts field is timestamp:incre. A cluster cannot query the oplog from mongos. timestamp can be converted into the Unix timestamp format based on the start time. The value of incre is 1.

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic .
Topic	Select the topic to be synchronized to the destination database.
Synchronize Topic To	The policy for synchronizing topics to the Kafka partitions. <ul style="list-style-type: none"> • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic.
Data Format in Kafka	Select the data format to be delivered to Kafka. <ul style="list-style-type: none"> • JSON: JSON message format, which is easy to interpret but takes up more space. For details, see Kafka Message Format .
Synchronization Object	The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-96 Task startup settings





Table 4-104 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.11 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 4-105 Supported databases

Source DB	Destination DB
RDS for PostgreSQL (versions 9.5, 9.6, 10, 11, 12, 13, 14 and 15)	<ul style="list-style-type: none"> On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15) ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15) <p>NOTE The major version of the destination database must be the same as or later than that of the source database.</p>

Supported Synchronization Objects

Table 4-106 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-106 Supported synchronization objects

Type	Notes
Objects	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases. ● Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. ● Scope of full synchronization <ul style="list-style-type: none"> - The following objects are supported in the database-level synchronization: schemas, tables, indexes, constraints, views, materialized views, sequences, stored procedures, rules, triggers, foreign keys, sorting rules, plug-ins, code conversion information, aggregate functions, operators, statistics extension, conversion information, text search configurations, functions, data types, type conversion, event triggers, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, and common indexes can be synchronized. During object file import, tables can be synchronized. - System objects such as system schemas, system tables, system users, and system functions cannot be synchronized. Other objects such as tablespaces, foreign data wrappers, foreign servers, user mappings, releases, subscriptions, and users cannot be synchronized. Object permissions cannot be synchronized. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain +"%'\<>, the schema name and table name cannot contain ".'\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. ● Function: Do not synchronize C language functions or functions with the leakproof or support attribute. ● Plug-in: The metadata of plug-ins is not synchronized. ● Data type: Basic data types are not synchronized. ● Type conversion: The binary coercion type cannot be converted. <ul style="list-style-type: none"> ● Scope of incremental synchronization <ul style="list-style-type: none"> - Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. - Some DDL statements can be synchronized, including TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA, CREATE TABLE, DROP TABLE, ALTER TABLE

Type	Notes
	<p>(including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW.</p> <p>During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, COMMENT ON COLUMN, COMMENT ON TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p> <ul style="list-style-type: none"> - Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE</p> <ul style="list-style-type: none"> • The source database captures DDL statements using event triggers and records them in specific tables, so you need to create event triggers and functions in the source database in advance. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL. • Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-107](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-107 Database account permission

Type	Full	Full+Incremental
Source database user	Database CONNECT permission, schema USAGE permission, table SELECT permission, and sequence SELECT permission	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permissions for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration.</p>
Destination database user	<ul style="list-style-type: none"> ● Database-level: <ul style="list-style-type: none"> - If the destination database is not PostgreSQL, the CREATEDB permission is required. - If the destination database is PostgreSQL, the CONNECT and CREATE permissions on PostgreSQL databases and the USAGE and CREATE permissions on public schemas are required. ● Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the object are required. <p>NOTE If the destination database version is PostgreSQL 15 and no database is created on the destination database for database-level synchronization and table-level synchronization, full or incremental synchronization may fail due to lack of the USAGE and CREATE permissions of the public schema.</p>	

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
 - Data-Level Comparison

To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. (A single incremental synchronization task or a single full synchronization task contains three phases.) To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-108 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The partition table trigger of the source database cannot be set to disable. - For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database. - To perform incremental synchronization: The wal_level value of the source database must be logical. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Source database object requirements: <ul style="list-style-type: none"> - Triggers with the same name cannot exist in the source database. - The objects that have dependencies must be synchronized at the same time. Otherwise, the synchronization may fail. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database cannot contain objects with the same type and name as the objects to be synchronized, including databases, schemas, and tables. System databases, system schemas, and system tables are excluded. - The destination table can contain more columns than the source table. However, the following failures must be avoided:

Type	Constraints
	<p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value +original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_+ hash value + original constraint name (which may be truncated) + _key. - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - For a full+incremental or incremental synchronization task, if an internal error occurs during the pre-check and the task stops before it is started, check and delete the streaming replication slot by referring to Forcibly Stopping Synchronization of PostgreSQL to avoid residual streaming replication slots in the source database. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the

Type	Constraints
	<p>source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later.</p> <ul style="list-style-type: none"> - If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database. <p>NOTE DDL statements are captured using event triggers in the source database, recorded in a specific table, and then synchronized to the destination database. You need to create event triggers and functions in the source database before starting a task. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, data may be inconsistent or the task may fail. • During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail.
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<ul style="list-style-type: none"> ● Stop a task normally. <ul style="list-style-type: none"> - The destination database sequence value is automatically reset. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. If users are synchronized, the user memberships are automatically synchronized after the task is complete. - When a full synchronization task is complete, the sequence values are compared based on logical consistency. That is, if an auto-increment sequence value is used only when the sequence value of the destination database is greater than or equal to that of the source database, or an auto-decrement sequence value is used only when the sequence value of the destination database is less than or equal to that of the source database, the sequence values in the source database are consistent with those in the destination database. - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. - If the value of session_replication_role of the destination database is replica when the full+incremental synchronization task is complete, change the value to the original one. ● Forcibly stop a task. <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is <code>drs_unique_ID</code>. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.
Troubleshooting	<ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Procedure

This section uses out-of-cloud synchronization from PostgreSQL to PostgreSQL as an example to describe how to configure a real-time synchronization task in the VPC network scenario.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-97 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.

The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region ⌵

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project ⌵

* Task Name ⓘ

DRS-5678

Description ⓘ

0/256

Table 4-109 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-98 Synchronization instance details

⚠ This information cannot be modified after you go to the next page.

Data Flow

To the cloud Out of the cloud Self Sync to Self Sync

Out of the cloud: The source database must be a database in the current cloud.

Source DB Engine

MySQL DDM GaussDB Distributed GaussDB Primary/Standby MariaDB DDS PostgreSQL GaussDB for MySQL

Destination DB Engine

Kafka PostgreSQL

Network Type

Public network ⓘ

✔ DRS will automatically bind the specified EP to the DRS instance and unbind the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

Source DB Instance

No DB instance available View DB Instance View Unavailable DB Instance

Synchronization Instance Name

Select the instance View Subnets View Occupied IP Address

Synchronization Mode

Full/Incremental Full Incremental

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization period logs to ensure data consistency between the source and destination databases.

Specify EP

 Create an EP

Table 4-110 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select PostgreSQL .
Network Type	<p>Available options: VPC, Public network, and VPN or Direct Connect. VPC is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The RDS for PostgreSQL DB instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-99 Task type




Table 4-111 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p>

- Enterprise Project and Tags

Figure 4-100 Enterprise Project and Tags




Table 4-112 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- Source database configuration

Figure 4-101 Source database information

Source Database

DB Instance Name

Database Username

Database Password

[Test Connection](#) This button is available only after the replication instance is created successfully.

Table 4-113 Source database settings

Parameter	Description
DB Instance Name	The RDS for PostgreSQL instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

- Destination database configuration

Figure 4-102 Destination database information

Destination Database

VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

[Test Connection](#) This button is available only after the replication instance is created successfully.

Table 4-114 Destination database settings

Parameter	Description
VPC	A dedicated virtual network in which the destination database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the destination database resides. You need to enable DHCP for creating the destination database subnet.
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-103 Synchronization Mode

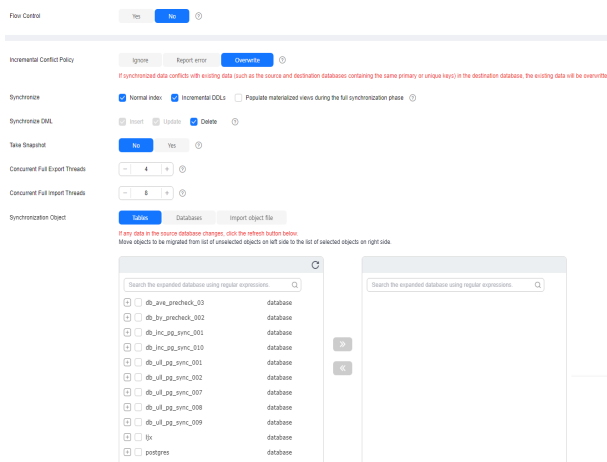
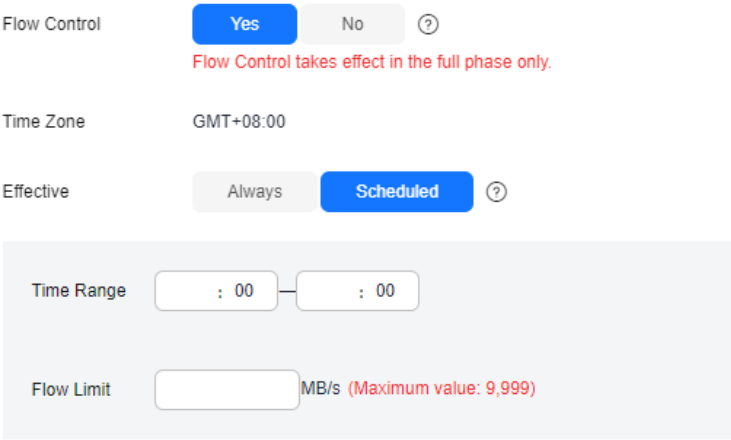



Table 4-115 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-104 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize	<p>Available options: Normal index, Incremental DDLs, and Populate materialized views during the full synchronization phase</p> <p>Populate materialized views during the full synchronization phase: This option takes effect only for materialized views that was populated in the source database. This operation affects the full synchronization performance. You perform this operation after the full synchronization is complete.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Take Snapshot	<p>If you perform a full+incremental migration, you can take a snapshot for your databases.</p> <p>Exporting data in snapshot mode in the full export phase can effectively improve the data synchronization efficiency in the full+incremental export scenario. However, the snapshot mechanism of PostgreSQL prevents historical data in the database from being reclaimed during the export, which may cause space expansion. You are advised to use this method when the full or incremental data volume is large and the source database disk space is sufficient.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • Database-level synchronization: In full synchronization, the selected databases and the inventory data of the database objects are synchronized. In incremental synchronization, the DML and some DDL statements of all tables except unlogged tables and temporary tables are synchronized. • Table-level synchronization: In full synchronization, the inventory data of the selected tables, sequences, views, or materialized views is synchronized. In incremental synchronization, the DML and some DDL statements of the selected tables are synchronized. • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, after the synchronization, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-105 Task startup settings

Table 4-116 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.12 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 4-117 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> RDS for PostgreSQL 9.5, 9.6, 10, 11, 12, 13, and 14 	<ul style="list-style-type: none"> Kafka 0.11 or later

Supported Synchronization Objects

Table 4-118 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-118 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases. Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. <p>NOTE The restrictions on synchronization object names are as follows: The database name cannot contain +"%\<>, the schema name and table name cannot contain ".!<>, and the column name cannot contain double quotation marks (") and single quotation marks (').</p> <ul style="list-style-type: none"> Scope of incremental synchronization <ul style="list-style-type: none"> Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. Not supported: DDL statements, DML statements of unlogged tables and temporary tables <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-119](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-119 Database account permission

Type	Incremental Synchronization
Source database user	The CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, and the permission to create replication connections

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-120 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> The wal_level value of the source database must be logical. The test_decoding plug-in has been installed on the source database. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - The source must be an RDS for PostgreSQL instance. - During database-level synchronization, incremental synchronization tasks cannot be edited. - During table-level synchronization, incremental synchronization tasks can be edited, but the database cannot be changed. - Before starting a synchronization task, ensure that no long transaction is started in the source database. Otherwise, the creation of the logical replication slot will be blocked, leading the task to fail. - After a task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not delete the primary key of the source database table. Otherwise, incremental data may be lost or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be lost or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, incremental data may be lost or the task may fail. • During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.

- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-106 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-121 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-107 Synchronization instance details

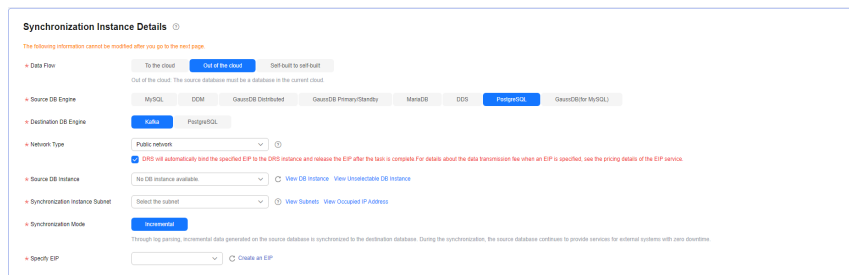


Table 4-122 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select Kafka .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Specifies an RDS PostgreSQL DB instance you have created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-108 Task type




Table 4-123 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p>

- Enterprise Project and Tags

Figure 4-109 Enterprise Project and Tags




Table 4-124 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-110 Source database information

Source Database

DB Instance Name

Database Username

Database Password

Table 4-125 Source database settings

Parameter	Description
DB Instance Name	The RDS for PostgreSQL instance selected when you created the migration task. The instance cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-111 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 4-126 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-112 Synchronization mode

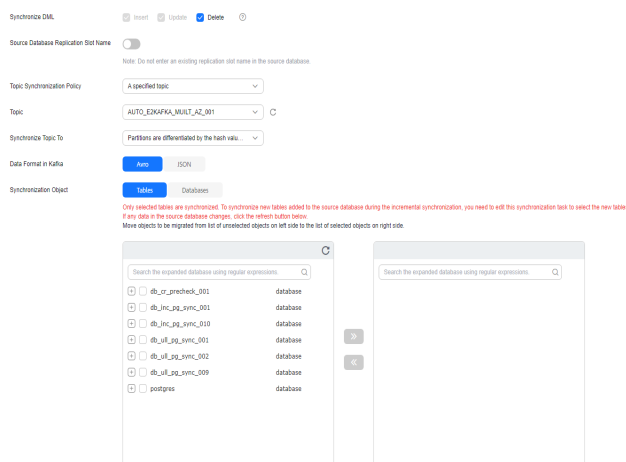



Table 4-127 Synchronization object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>

Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database, schema and table names, the performance on a single table query can be improved. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash value of the database_name.schema.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Data Format in Kafka	<p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro: A binary encoded format that is efficient. You need to deserialize the data later. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-113 Task startup settings

Table 4-128 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.13 From GaussDB Distributed to MySQL

Supported Source and Destination Databases

Table 4-129 Supported databases

Source DB	Destination DB
GaussDB distributed	<ul style="list-style-type: none">• RDS for MySQL 5.6 and 5.7• On-premises MySQL 5.5, 5.6, and 5.7 databases• MySQL 5.5, 5.6, and 5.7 databases on an ECS• MySQL 5.5, 5.6, and 5.7 databases on other clouds

Supported Synchronization Objects

Table 4-130 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-130 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-131](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-131 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user must have the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		details, see Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES		

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
 - To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-132 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the UTF8MB4 character set. If the source database uses the UTF8 character set and the destination database uses the UTF8MB3 character set, or if the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to inconsistent character set encoding ranges (GB18030_2022 > GB18030 > GBK > GB2312) or database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● The destination database object must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before synchronization, ensure that the table structure of the destination database has been created and is the same as that of the source database or contains all columns in the source database. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Restrictions
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - By default, the GaussDB-to-MySQL synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - The names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: select working_version_num();

Type	Restrictions
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.

Type	Restrictions
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not limit the synchronization speed during data comparison. MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-114 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0256

Table 4-133 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-115 Synchronization instance details

Synchronization Instance Details ⓘ

⚠ The following information cannot be modified after you go to the next page.

- Data Flow: To the cloud Out of the cloud Self-built to self-built
- Source DB Engine: MySQL, ODM, GaussDB(DWS), GaussDB Primary/Standby, MariaDB, OSS, PostgreSQL, GaussDB(for MySQL)
- Destination DB Engine: MySQL, Oracle, GaussDB(DWS), GaussDB Distributed, GaussDB Primary/Standby, Kafka
- Network Type: Public network (selected), Private network
- Source DB Instance: No DB instance available (selected), View DB Instance, View Unavailable DB Instance
- Synchronization Instance Subnet: Select the subnet (selected), View Subnets
- Synchronization Mode: Full (selected), Incremental
- Source DB Quantity: 2
- Specify EIP: Create an EIP

This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Table 4-134 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select MySQL .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Incremental, and Full</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-116 Task type




Table 4-135 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-117 Enterprise Project and Tags




Table 4-136 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-118 Source database information

Source Database

DB Instance Name

Database Username


Database Password 

Table 4-137 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-119 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 4-138 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-120 Synchronization mode

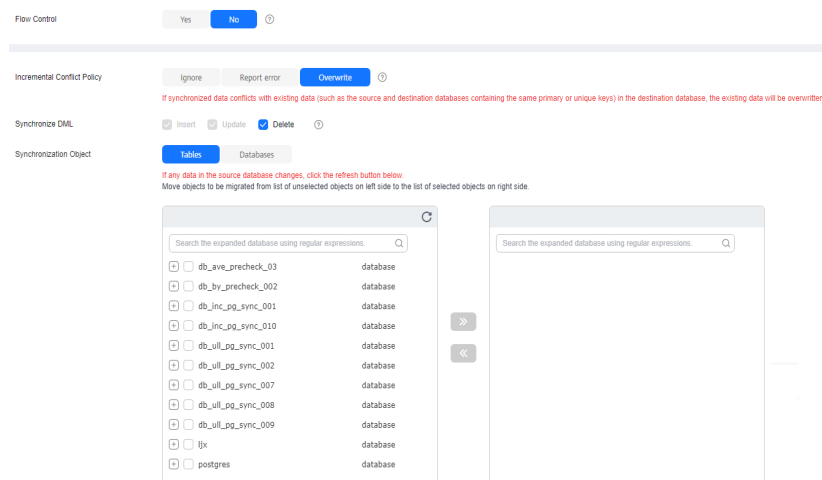
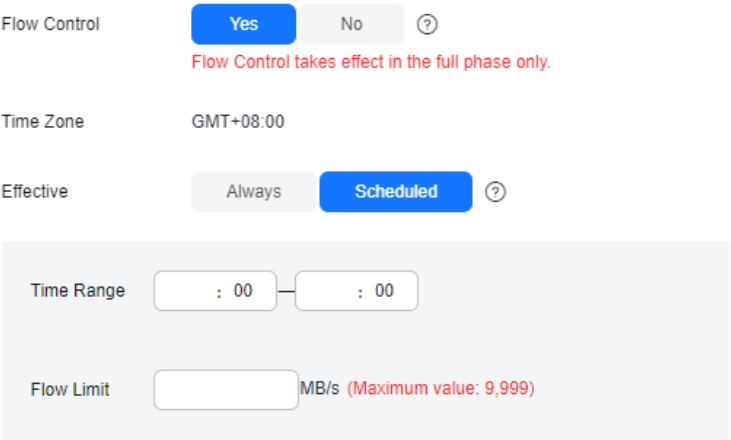



Table 4-139 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-121 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-122 Task startup settings

Table 4-140 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.14 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 4-141 Supported databases

Source DB	Destination DB
GaussDB distributed	<ul style="list-style-type: none">• On-premises Oracle 11g and 19c databases• Oracle 11g and 19c databases on an ECS

Supported Synchronization Objects

Table 4-142 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-142 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\` \?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-143](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-143 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		<p>Connection and Port Description for Incremental Synchronization from GaussDB.</p>	<p>a GaussDB Database.</p> <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role.</p>	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</p>	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</p>

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-144 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the AL32UTF8 or UTF8 character set. If the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the ZHS16GBK, ZHS16CGB231280, or ZHS32GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● Destination database object requirements: <ul style="list-style-type: none"> - Before synchronization, ensure that the corresponding table structure has been created in the destination database and is the same as that in the source instance. If column processing is performed, ensure the consistency of the table structure after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination database has sufficient disk space. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.

Type	Restrictions
	<p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - Synchronization from GaussDB distributed to Oracle is the backward synchronization process for synchronization from Oracle to GaussDB distributed. You are not advised to use either of the two synchronization modes independently. - By default, the GaussDB-to-Oracle synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The selected source database cannot contain tables with the same name but different cases. Otherwise, the synchronization fails. You are advised to synchronize only the schema and table names that are in uppercase on the Oracle database, and the schema and table names that are in lowercase on the GaussDB distributed database. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in

Type	Restrictions
	<p>a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails.</p> <ul style="list-style-type: none"> - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● For tables containing large fields (including blob, clob, nclob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. ● If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison.

Type	Restrictions
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-123 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: ⓘ

Description: ⓘ

0/256

Table 4-145 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-124 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

* Data Flow: To the cloud Out of the cloud Self-built to self-built

* Source DB Engine: MySQL, ODM, **GaussDB Distributed**, GaussDB Primary/Standby, MariaDB, DDS, PostgreSQL, GaussDB(for MySQL)

* Destination DB Engine: MySQL, **GaussDB(DWS)**, GaussDB Distributed, GaussDB Primary/Standby, Kafka

* Network Type: Public network

* Source DB Instance: No DB instance available.

* Synchronization Instance Subnet: Select the subnet.

* Synchronization Mode: **Full/Incremental**, Full, Incremental

* Source DB Quantity:

* Specify EP:

Table 4-146 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .

Parameter	Description
Destination DB Engine	Select Oracle .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-125 Task type




Table 4-147 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-126 Enterprise Project and Tags




Table 4-148 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-127 Source database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 4-149 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-128 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

Database Username

Database Password

SSL Connection

Table 4-150 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance.
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-129 Synchronization mode

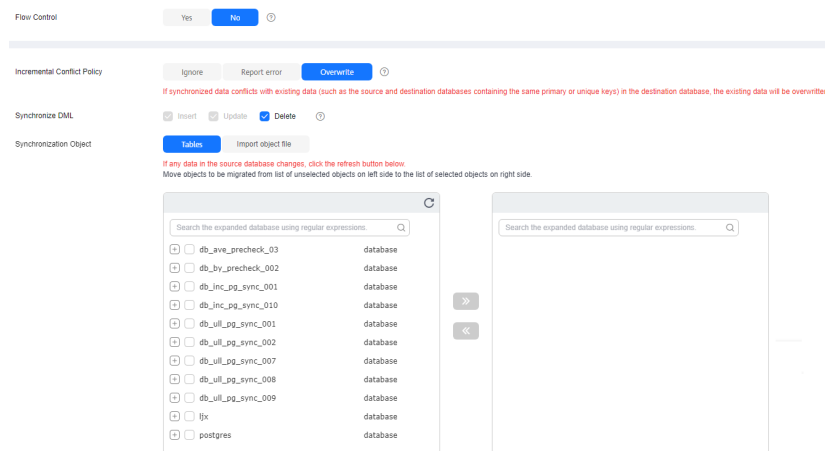
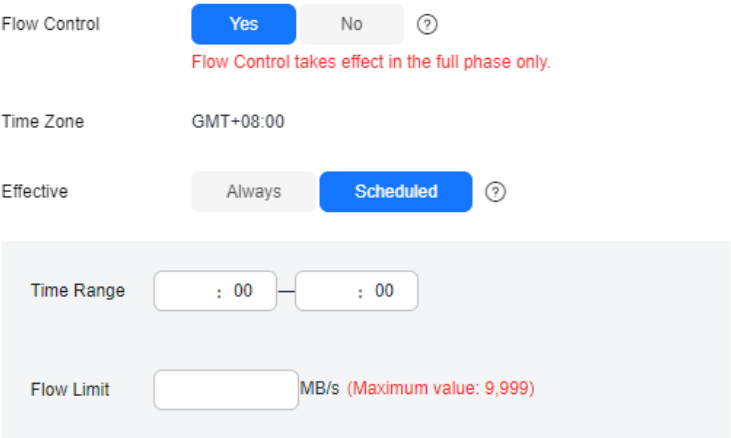



Table 4-151 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-130 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-131 Task startup settings

Table 4-152 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.15 From GaussDB Distributed to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-153 Supported databases

Source DB	Destination DB
GaussDB distributed	GaussDB(DWS) cluster

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-154 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-154 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY • Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only table data, table structures, and index constraints can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Synchronization Scope
	<ul style="list-style-type: none"> - Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. - When many-to-one mapping is required, you need to create a table structure in the destination database that is the same as the table structure in the source database of the many-to-one synchronization.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-155](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-155 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, and SELECT or UPDATE (required only for tables without primary keys) permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	The user must have the sysadmin role or the following permissions: <ul style="list-style-type: none"> The CONNECT and CREATE permissions on databases The USAGE permission on schemas The INSERT, DELETE, UPDATE, and SELECT permissions on tables 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-156 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` ?!. - The distribution key of the source database cannot contain commas (,) and spaces. - The primary key and unique constraint of the source database table must contain the distribution column. Otherwise, the table structure cannot be synchronized. You need to manually create the table structure in the destination database and retry the task. - Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before executing an incremental synchronization task, create schemas and tables in the destination database. - For an incremental synchronization task, disable foreign keys for tables in the destination database to ensure that the sequence in which data is written to the destination database is the same as that in the source database. Otherwise, the synchronization task may fail. - The destination table can contain more columns than the source table. However, the following failures must be avoided:

Type	Constraints
	<p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>

Type	Constraints
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-132 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: ⓘ

Description: ⓘ 0/256

Table 4-157 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-133 Synchronization instance details

Synchronization Instance Details ⓘ

⚠ The following information cannot be modified after you go to the next page.

Data Flow: To the cloud **Out of the cloud** Self-host to self-host Out of the cloud: The source database resides in a database in the current cloud.

Source DB Engine: MySQL ODM **GaussDB Distributed** GaussDB Primary/Standby MariaDB DCS PostgreSQL GaussDB for MySQL

Destination DB Engine: MySQL Oracle **GaussDB Distributed** GaussDB Primary/Standby Kafka

Network Type: Public network Private network DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

Source DB Instance: No DB instance available View DB instance View Unavailable DB instance

Synchronization Instance Subnet: Select the subnet View Subnets

Synchronization Mode: **Full incremental** Full Incremental This synchronization task synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Source CN Quantity:

Specify EP: ⓘ

Table 4-158 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .

Parameter	Description
Destination DB Engine	Select GaussDB(DWS) .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-134 Task type




Table 4-159 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-135 Enterprise Project and Tags




Table 4-160 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-136 Source database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 4-161 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-137 Destination database information

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

Table 4-162 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-138 Synchronization mode

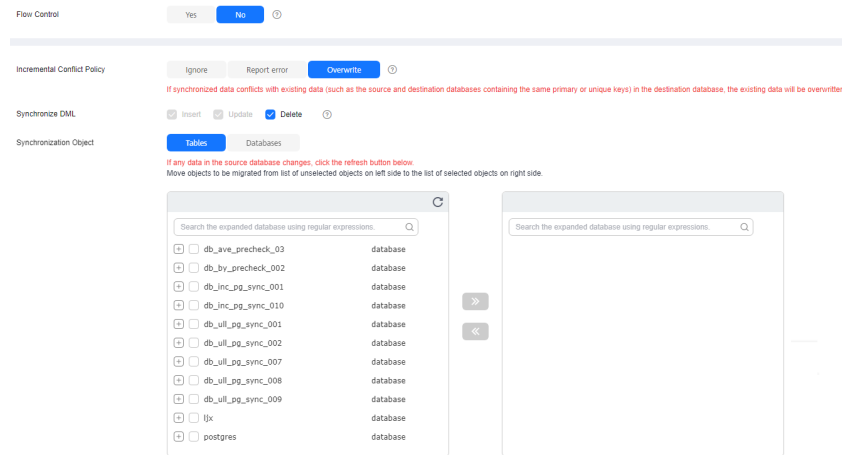
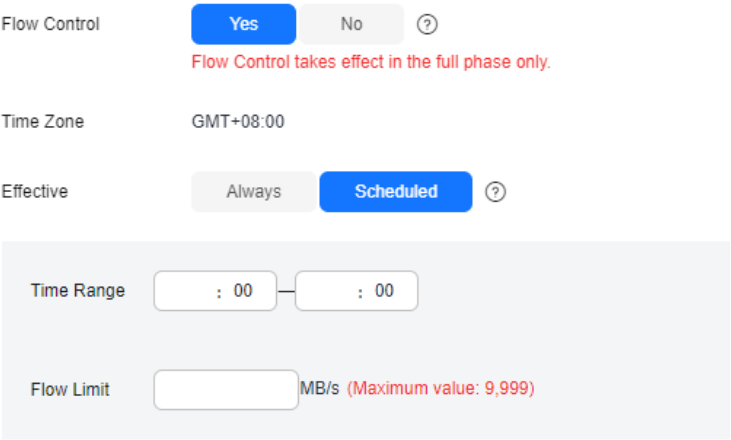



Table 4-163 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-139 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-140 Task startup settings

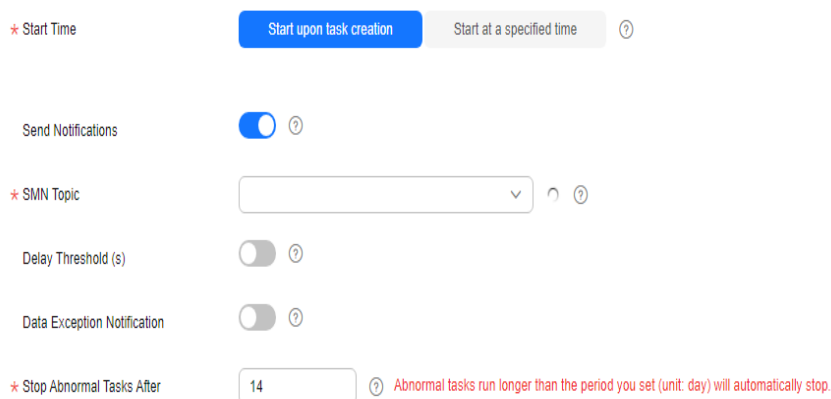



Table 4-164 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.16 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 4-165 Supported databases

Source DB	Destination DB
GaussDB distributed	Kafka 0.11 or later

Supported Synchronization Objects

Table 4-166 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-166 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenario: Incremental synchronization • Supported fields: BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and TINYINT • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of the selected table can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none"> The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-167](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-167 Database user permission

Type	Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a

synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-168 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` \?! ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME

Type	Constraints
	<p>ZONE data cannot be synchronized. To query the source database version, run the following SQL statement:</p> <pre>select working_version_num();</pre>
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the REPLICIA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● During migration of table-level objects, you are not advised to rename the tables. ● Replication of interval partition tables is not supported. ● The name of a primary key column cannot be changed. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-141 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: ⓘ

Description: ⓘ 0/256

Table 4-169 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-142 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud | [Use of the cloud](#) | [Switch to self-host](#)
Out of the cloud. The source database must be a database in the current cloud.

Source DB Engine: MySQL, DDM, **GaussDB Distributed**, GaussDB Primary/Standby, MariaDB, DDS, PostgreSQL, GaussDB(for MySQL)

Destination DB Engine: MySQL, Oracle, GaussDB(DWS), GaussDB Distributed, GaussDB Primary/Standby, **Auto**

Network Type: Public network | ⓘ
DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is completed. For details about the data transmission fee when an EIP is specified, see the pricing details of the DRS service.

DRS Task Type: [Simple DRS](#) | [Standard](#)
Instances deployed in a user's account. The synchronization task will be created on only one side to save costs. This deployment is for scenarios where there is a small amount of source data, instant synchronization is required, and there is no requirement in source databases.

Source DB Instance: No DB instance available | ⓘ [View DB Instance](#) | [View Unavailable DB Instance](#)

Synchronization Instance Subnet: Select the subnet | ⓘ [View Subnets](#)

Synchronization Mode: [Incremental](#)
Through log parsing, incremental data generated on the source database is synchronized to the destination database. During the synchronization, the source database continues to provide services for external systems with zero downtime.

Source DB Quantity: 1 | ⓘ

Specify EIP: [Dropdown menu] | ⓘ [Create an EIP](#)

Table 4-170 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .

Parameter	Description
Destination DB Engine	Select Kafka .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The GaussDB distributed instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-143 Task type




Table 4-171 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-144 Enterprise Project and Tags




Table 4-172 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-145 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-173 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-146 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

● Test successful

Table 4-174 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-147 Synchronization mode

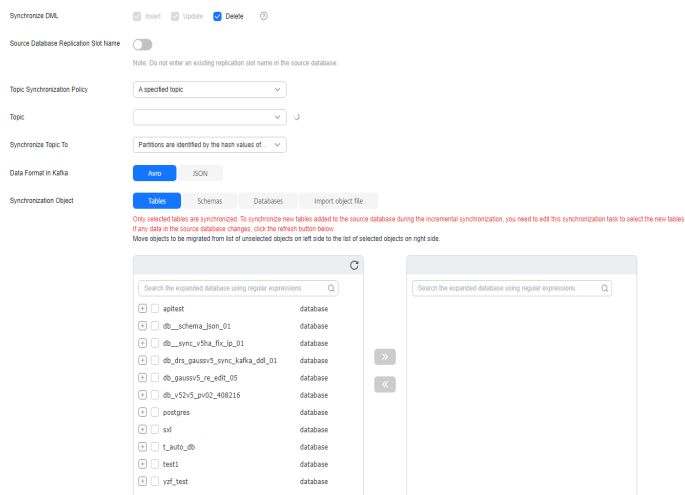



Table 4-175 Synchronization Object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .

Parameter	Description
Topic Name Format	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database, schema and table names, the performance on a single table query can be improved. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash value of the database_name.schema.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.

Parameter	Description
Data Format in Kafka	<p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> ● Avro refers to binary encoded format. ● JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-148 Task startup settings





Table 4-176 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.17 From GaussDB Distributed to GaussDB Distributed

Supported Source and Destination Databases

Table 4-177 Supported databases

Source DB	Destination DB
GaussDB distributed	<p>GaussDB distributed</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p>

Supported Synchronization Objects

Table 4-178 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-178 Supported synchronization objects

Type	Constraints
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none">- Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.- The database name, schema name, and table name cannot contain special characters /<.>\\' ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-179](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-179 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB .	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • To use value comparison, you must have the MONADMIN permission. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-180 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Constraints
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - GaussDB distributed supports only ranged partitioned tables. Therefore, other types of partitioned tables are synchronized to the destination database as common tables. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. The topology structure of the source database must be the same as that of the destination database. Otherwise, the content comparison function is unavailable. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-149 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 4-181 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-150 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

- **Data Flow**
 - To the cloud **Out of the cloud** Self-built to self-built
 - Out of the cloud: The source database must be a database in the current cloud.
- **Source DB Engine**
 - MySQL CDM **GaussDB Distributed** GaussDB Primary/Standby MariaDB DDS PostgreSQL GaussDB for MySQL
- **Destination DB Engine**
 - MySQL Oracle GaussDB(DRS) **GaussDB Distributed** GaussDB Primary/Standby Fuba
- **Network Type**
 - Public network Private network
 - DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.
- **Source DB Instance**
 - No DB instance available
 - [View DB Instance](#) [View Unavailable DB Instance](#)
- **Synchronization Instance Subnet**
 - Select the subnet
 - [View Subnets](#)
- **Task Mode**
 - Single **Multiple**
 - In single-task mode, data is extracted from the CK. In multi-task mode, data is extracted from DMS. Each data node corresponds to a single address.
- **Synchronization Mode**
 - Full Synchronization** Full Incremental
 - This synchronization type synchronizes data in real time, after a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
- **Source CK Quantity**
 -
- **Specify EIP**
 -
 - [Create an EIP](#)

Table 4-182 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Task Mode	<p>The options are Single and Multiple.</p> <p>In single-task mode, data is extracted from the CN. In multi-task mode, data is extracted from DNs. Each data node corresponds to a single subtask.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-151 Task type




Table 4-183 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-152 Enterprise Project and Tags




Table 4-184 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-153 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-185 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-154 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 4-186 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-155 Synchronization mode

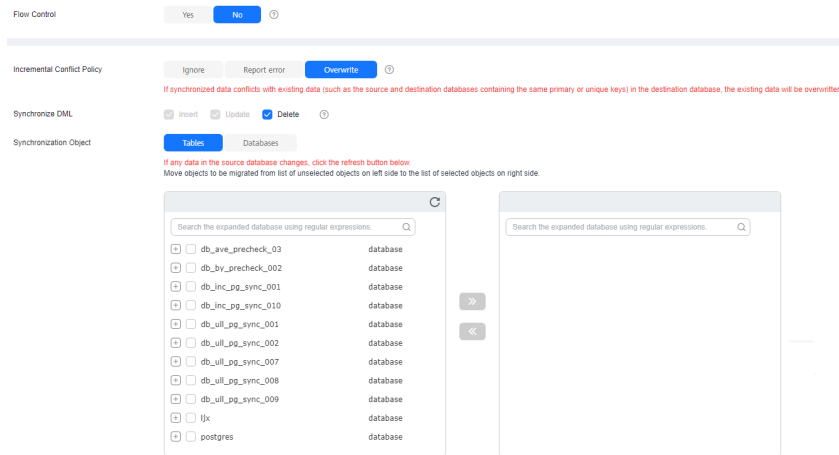
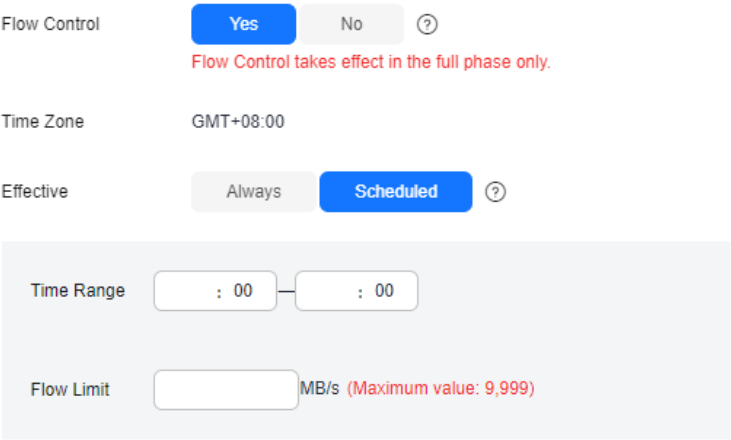



Table 4-187 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-156 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-157 Task startup settings

Table 4-188 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.18 From GaussDB Distributed to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-189 Supported databases

Source DB	Destination DB
GaussDB distributed	GaussDB primary/standby NOTE The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

Table 4-190 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-190 Supported synchronization objects

Type	Constraints
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none">- Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.- The database name, schema name, and table name cannot contain special characters /<.>\\' ! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-191](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-191 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		<p>Connection and Port Description for Incremental Synchronization from GaussDB.</p>	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-192 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Constraints
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - GaussDB distributed supports only ranged partitioned tables. Therefore, other types of partitioned tables are synchronized to the destination database as common tables. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-158 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Project

* Task Name

Description

0/256

Table 4-193 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-159 Synchronization instance details

Synchronization Instance Details

Data Flow To the cloud Out of the cloud Self-host to self-host

Out of the cloud: The source database must be a database in the current cloud.

Source DB Engine: MySQL, EDW, GaussDB Database, GaussDB Primary/Standby, MariaDB, DDS, PostgreSQL, GaussDBfor MySQL

Destination DB Engine: MySQL, Oracle, GaussDB(DWS), GaussDB Distributed, GaussDB Primary/Standby, Kafka

Network Type: Public network

Source DB Instance: No DB instance available, New DB Instance, View Unavailable DB Instance

Synchronization Instance Subnet: Select the subnet, New Subnet

Synchronization Mode: Full/Incremental, Full, Incremental

Source DB Quantity: 1, 2, 3, 4

Security EP: Create an EP

Table 4-194 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-160 Task type




Table 4-195 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-161 Enterprise Project and Tags




Table 4-196 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-162 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-197 Source database settings

Parameter	Description
DB Instance Name	The GaussDB distributed DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-163 Destination database information

Destination Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

Test Connection This button is available only after the replication instance is created successfully.

Table 4-198 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-164 Synchronization mode

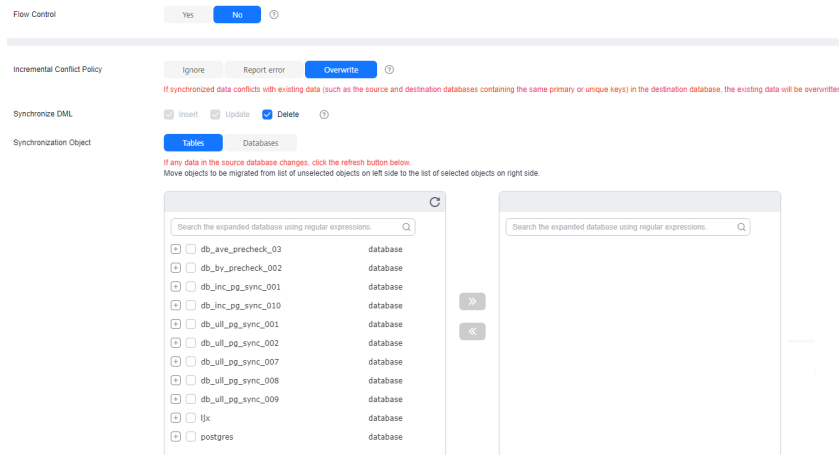
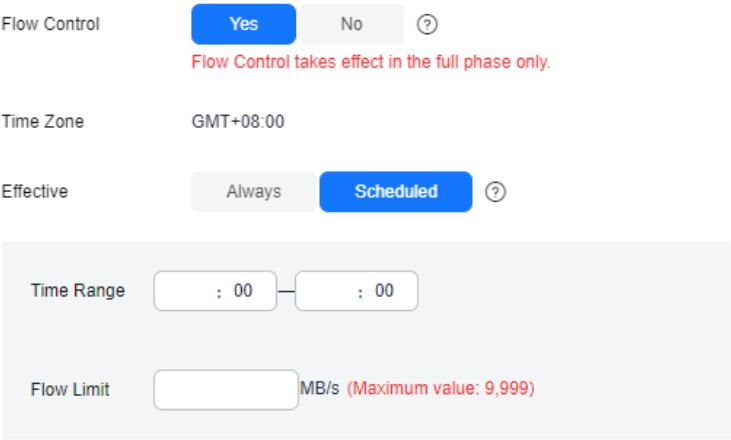



Table 4-199 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-165 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process.● Overwrite Conflicting data will be overwritten.● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-166 Task startup settings

Table 4-200 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.19 From GaussDB Distributed to PostgreSQL

Supported Source and Destination Databases

Table 4-201 Supported databases

Source DB	Destination DB
GaussDB Distributed	<ul style="list-style-type: none">• On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)• ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-202 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-202 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● The following field types are not supported: SMALLDATETIME, REETIME, ABSTIME, TID, XID, CID and OID ● Table-level synchronization is supported. <ul style="list-style-type: none"> - The table data can be synchronized. - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-203](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-203 Database user permissions

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	The user must have the following minimum permissions: <ul style="list-style-type: none"> • Database permission: CONNECT • Schema permission: USAGE • Table permission: INSERT, UPDATE, DELETE, and SELECT • Sequence permission: UPDATE 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-204 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The destination database, schemas, and table object structures must be created in advance. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Constraints
	<ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● When GaussDB is not compatible with PostgreSQL, the processing precision of the date data type in the GaussDB database may be different from that in the PostgreSQL database. As a result, data of the two databases is inconsistent.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-167 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region .. v

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-205 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-168 Synchronization instance details

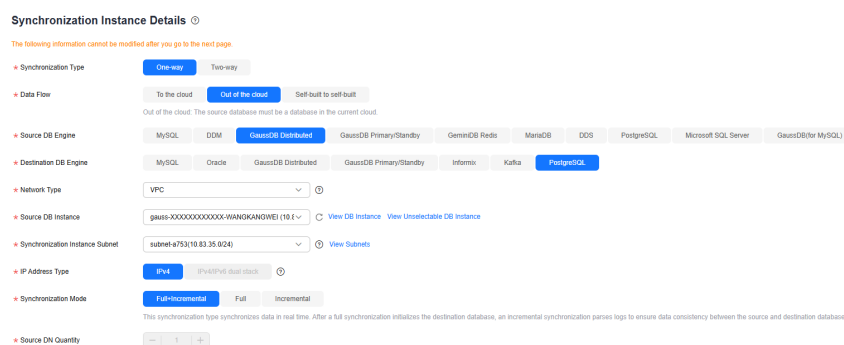


Table 4-206 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select PostgreSQL .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The distributed GaussDB instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.

Parameter	Description
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-169 AZ



Table 4-207 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-170 Enterprise Project and Tags




Table 4-208 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-171 Source database information

Source Database

DB Instance Name gauss-XXXXXXXXXX-WANGKANGWEI (10.83.35.16:8000)

Database Username

Database Password

Test Connection This button is available only after the replication instance is created successfully.

Table 4-209 Source database settings

Parameter	Description
DB Instance Name	The distributed GaussDB DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-172 Destination database information

Destination Database

Select Connection

VPC vpc-3962172-19-0-012 [View VPC](#)

Subnet subnet-2669172-19-0-024 [View Subnets](#)

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the destination database, and related parameters have been correctly configured.

Test Connection This button is available only after the replication instance is created successfully.

Table 4-210 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-173 Synchronization Mode

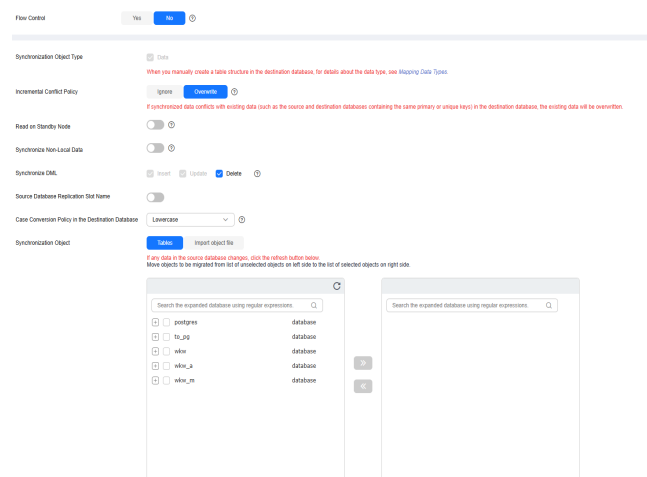
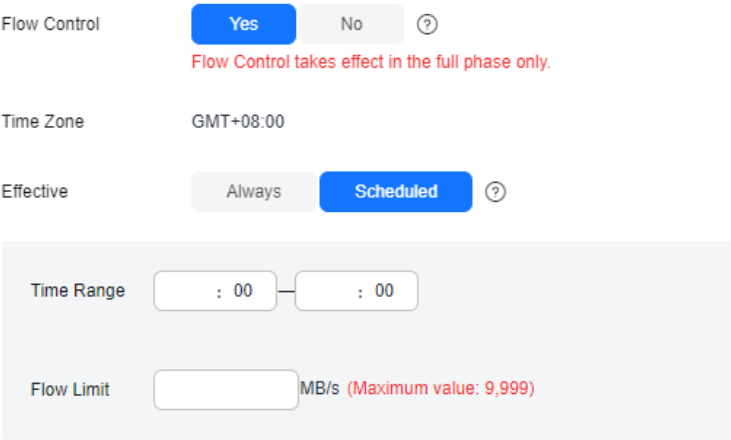



Table 4-211 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-174 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	Type of objects for full synchronization. Data is mandatory.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.
Case Conversion Policy in the Destination Database	Case of schema names, table names, and column names can be converted. If you have specified the mapping name when selecting objects to be synchronized, ignore this parameter.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-175 Task startup settings





Table 4-212 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.20 From GaussDB Primary/Standby to MySQL

Supported Source and Destination Databases

Table 4-213 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<ul style="list-style-type: none"> RDS for MySQL 5.6 and 5.7 On-premises MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on an ECS MySQL 5.5, 5.6, and 5.7 databases on other clouds

Supported Synchronization Objects

Table 4-214 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-214 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-215](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-215 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-216 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● The destination database parameters must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The source database supports the UTF8 character set, and the destination database supports the UTF8MB4 character set. If the source database uses the UTF8 character set and the destination database uses the UTF8MB3 character set, or if the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to inconsistent character set encoding ranges (GB18030_2022 > GB18030 > GBK > GB2312) or database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● The destination database object must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before synchronization, ensure that the table structure of the destination database has been created and is the same as that of the source database or contains all columns in the source database. - The destination table can contain more columns than the source table. However, the following failures must be avoided:

Type	Restrictions
	<p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - By default, the GaussDB-to-MySQL synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - In GaussDB, if there is 0x00 in character or binary data, the data will be truncated, which is inconsistent with MySQL. As a result, the final data may be inconsistent.

Type	Restrictions
	<ul style="list-style-type: none"> - Binary data cannot be used as a DML matching condition. That is, binary data cannot be used as a primary key. Binary data is fixed in length and is automatically filled in. The fill character of GaussDB is 0x20, and that of MySQL is 0x00. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent.

Type	Restrictions
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-176 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name:

Description:

0/256

Table 4-217 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-177 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud | **Out of the cloud** | Self-built to self-built

Source DB Engine: MySQL | DDM | GaussDB Distributed | **GaussDB Primary/Standby** | GaussDB | DDS | PostgreSQL | GaussDB for MySQL

Destination DB Engine: MySQL | Oracle | GaussDB(DWS) | GaussDB Distributed | GaussDB Primary/Standby | Kafka

Network Type:

Source DB Instance: No DB Instance Available | View DB Instance | View Unavailable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnets | View Occupied IP Address

Synchronization Mode: Full-Incremental | Full | Incremental

Specify EP: | Create an EP

Table 4-218 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select MySQL .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-178 Task type




Table 4-219 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-179 Enterprise Project and Tags




Table 4-220 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-180 Source database information

Source Database

DB Instance Name

Database Username

Database Password

Table 4-221 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-181 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-222 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-182 Synchronization mode

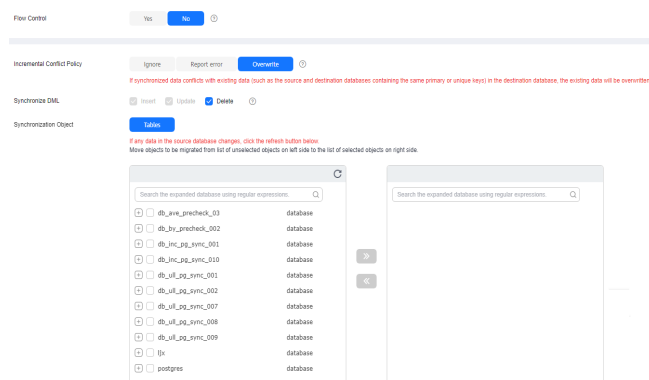
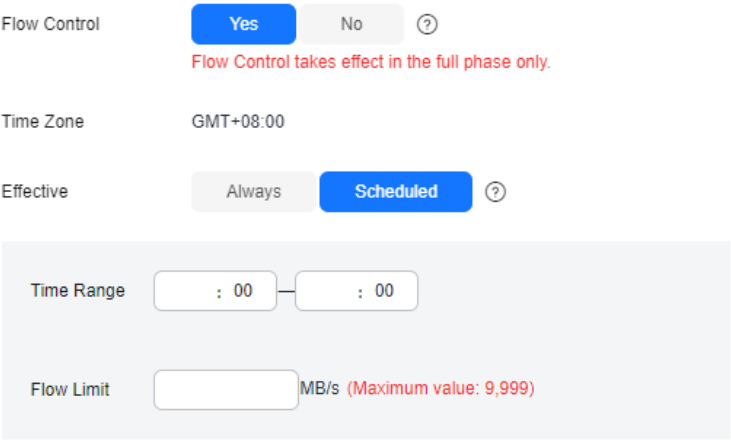



Table 4-223 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-183 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names).</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-184 Task startup settings

Table 4-224 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.21 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 4-225 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<ul style="list-style-type: none">• On-premises Oracle 11g and 19c databases• Oracle 11g and 19c databases on an ECS

Supported Synchronization Objects

Table 4-226 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-226 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML and some DDL statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.).

Type	Synchronization Scope
	<ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-227](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-227 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role.	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-228 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the AL32UTF8 or UTF8 character set. If the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the ZHS16GBK, ZHS16CGB231280, or ZHS32GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● Destination database object requirements: <ul style="list-style-type: none"> - Before synchronization, ensure that the corresponding table structure has been created in the destination database and is the same as that in the source instance. If column processing is performed, ensure the consistency of the table structure after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination database has sufficient disk space. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.

Type	Restrictions
	<p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - Synchronization from GaussDB primary/standby to Oracle is the backward synchronization process for synchronization from Oracle to GaussDB primary/standby. You are not advised to use either of the two synchronization modes independently. - By default, the GaussDB-to-Oracle synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The selected source database cannot contain tables with the same name but different cases. Otherwise, the synchronization fails. You are advised to synchronize only the schema and table names that are in uppercase on the Oracle database, and the schema and table names that are in lowercase on the GaussDB primary/standby database. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.

Type	Restrictions
	<ul style="list-style-type: none"> - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● For tables containing large fields (including blob, clob, nlob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. ● If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type. ● The following types of DDL statements cannot be synchronized: ALTER TABLE MERGE PARTITIONS and CREATE INDEX CONCURRENTLY. ● To synchronize DDL statements, ensure that the compatibility modes of the source and destination databases are the same. ● For DDL statements involving tablespace operations, the user of the destination database must have the tablespace operation permissions.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out.

Type	Restrictions
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-185 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: [Text input: DRS-5678] ⓘ

Description: [Text area: 0/256] ⓘ

Table 4-229 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-186 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow: **To the cloud** (Selected) | Out of the cloud | Default to inherit

• Source DB Engine: MySQL | OceanBase | GaussDB (Doris) | GaussDB (Distributed) | GaussDB (Primary/Standby) | MaxCompute | DWS | PostgreSQL | GaussDB (Elastic MySQL)

• Destination DB Engine: MySQL | OceanBase | GaussDB (Doris) | GaussDB (Distributed) | GaussDB (Primary/Standby) | Kafka

• Network Type: Public network | VPC

• Source DB Instance: [Dropdown menu] DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

• Synchronization Instance Subnet: [Dropdown menu] Select the subnet. | New Subnet | View Occupied IP Address

• Synchronization Mode: Full | Incremental

• Specify EP: [Dropdown menu] This synchronization task synchronizes data in real time. After a full synchronization releases the destination database, an incremental synchronization jumps logs to ensure data consistency between the source and destination databases. | Create an EP

Table 4-230 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .

Parameter	Description
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select Oracle .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-187 Task type




Table 4-231 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-188 Enterprise Project and Tags




Table 4-232 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-189 Source database information

Source Database

DB Instance Name

Database Username


Database Password 

Table 4-233 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-190 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

Database Username

Database Password

SSL Connection

Table 4-234 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance.
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-191 Synchronization mode

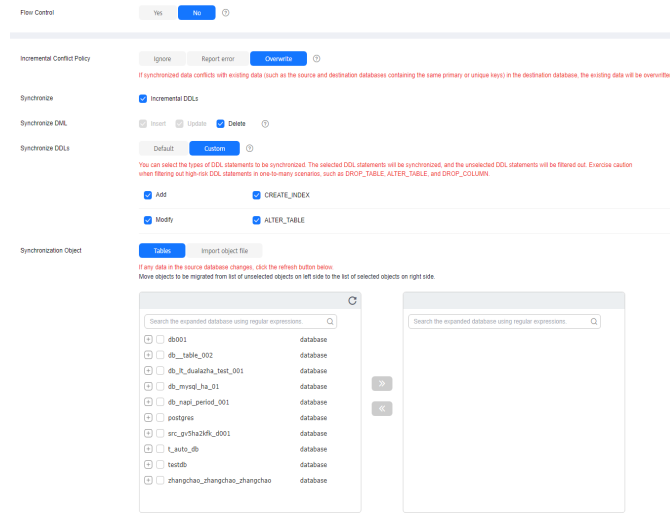
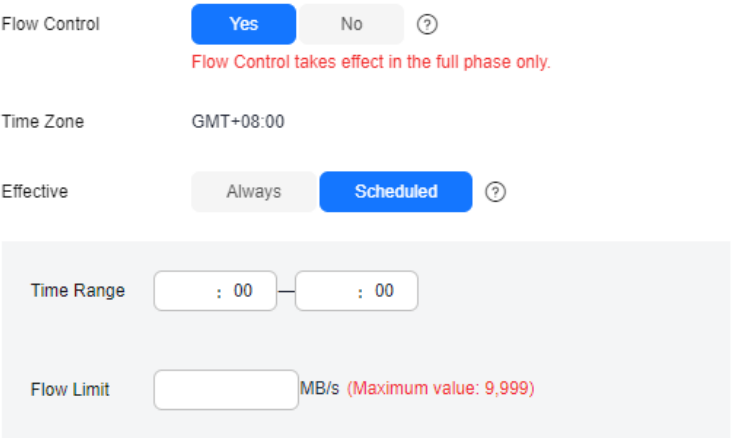



Table 4-235 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-192 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize	Indicates whether to synchronize DDLs in the incremental synchronization phase.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default: DDL statements of the selected synchronization objects will be synchronized to the destination database. <ul style="list-style-type: none"> - DDL statements supported by table-level synchronization: ALTER TABLE and CREATE INDEX ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. Only selected DDL types can be synchronized.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-193 Task startup settings

Table 4-236 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.22 From GaussDB Primary/Standby to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-237 Supported databases

Source DB	Destination DB
GaussDB primary/standby	GaussDB(DWS) cluster

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-238 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-238 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only table data, table structures, and index constraints can be synchronized. - During incremental synchronization, DRS allows you to synchronize DML statements of selected tables only. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbperf, pkg_service, pkg_util, db_file, db_random, db_output, db_raw, db_sql, db_lob, db_task, blockchain, db4ai, db_pldebugger, sqladvisor, db_application_info, db_match, db_pldeveloper, db_scheduler, db_session, db_utility, db_sql_util, db_xml, db_xmldom, db_xmlparser, db_compression, db_heat_map, db_ilm, db_ilm_admin, prvt_ilm, db_profiler, db_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none"> - The database name, schema name, and table name cannot contain special characters /<.>\\` ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. - If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. - When many-to-one mapping is required, you need to create a table structure in the destination database that is the same as the table structure in the source database of the many-to-one synchronization.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-239](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-239 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, and SELECT or UPDATE (required only for tables without primary keys) permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		<p>Connection and Port Description for Incremental Synchronization from GaussDB.</p>	<p>a GaussDB Database.</p> <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	<p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT and CREATE permissions on databases The USAGE and CREATE permissions on schemas The INSERT, DELETE, UPDATE, and SELECT permissions on tables 	<p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT permission on databases The USAGE permission on schemas The INSERT, DELETE, UPDATE, and SELECT permissions on tables 	<p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT and CREATE permissions on databases The USAGE and CREATE permissions on schemas The INSERT, DELETE, UPDATE, and SELECT permissions on tables

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"?! - The distribution key of the source database cannot contain commas (,) and spaces. - Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before executing an incremental synchronization task, create schemas and tables in the destination database. - For an incremental synchronization task, disable foreign keys for tables in the destination database to ensure that the sequence in which data is written to the destination database is the same as that in the source database. Otherwise, the synchronization task may fail. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.

Type	Constraints
	<p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - When a schema name or table name is mapped, to prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: <code>i_+hash value+original index name (which may be truncated)+_key</code>. The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: <code>c_+hash value+original constraint name (which may be truncated)+_key</code>. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), <code>TIMESTAMP WITH TIME ZONE</code> data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>

Type	Constraints
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-194 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 4-240 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-195 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

* Data Flow: To the cloud **Out of the cloud** Self-built to self-built

Out of the cloud: The source database must be a database in the current cloud.

* Source DB Engine: MySQL, ODM, GaussDB Distributed, **GaussDB Primary/Standby**, GaussDB, DDS, PostgreSQL, GaussDB for MySQL

* Destination DB Engine: MySQL, Oracle, **Advanced Oracle**, GaussDB Distributed, GaussDB Primary/Standby, Kafka

* Network Type: Public network Private network

DRS will automatically bind the specified EP to the DRS instance and release the EP after the test is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the DRS service.

* Source DB Instance:

* Synchronization Instance Subnet:

* Synchronization Mode: Full Incremental Full Incremental

This synchronization type synchronizes data in real time, offers a full synchronization initiation, an incremental synchronization pattern to ensure data consistency between the source and destination databases.

* Standby EP:

Table 4-241 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB(DWS) .

Parameter	Description
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-196 Task type




Table 4-242 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-197 Enterprise Project and Tags




Table 4-243 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-198 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-244 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-199 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

Table 4-245 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-200 Synchronization mode

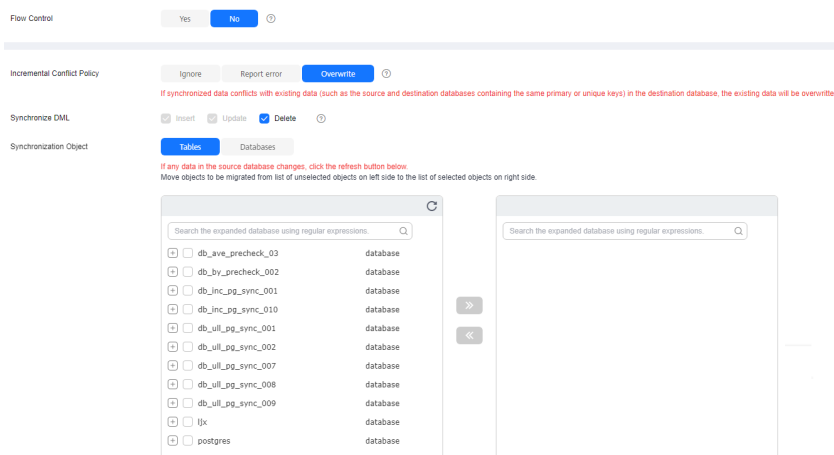
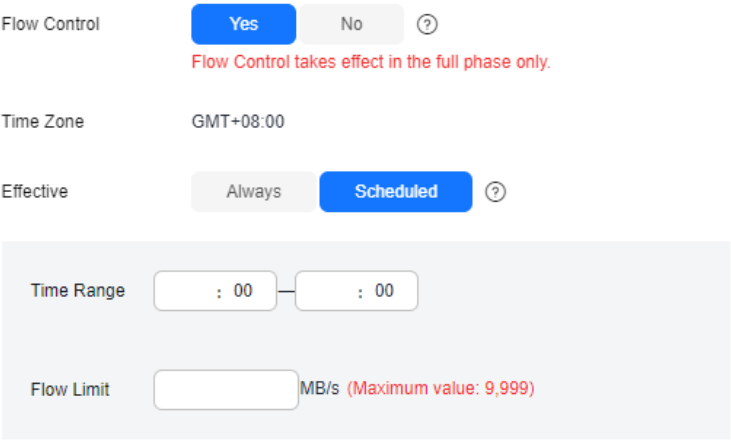



Table 4-246 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-201 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-202 Task startup settings

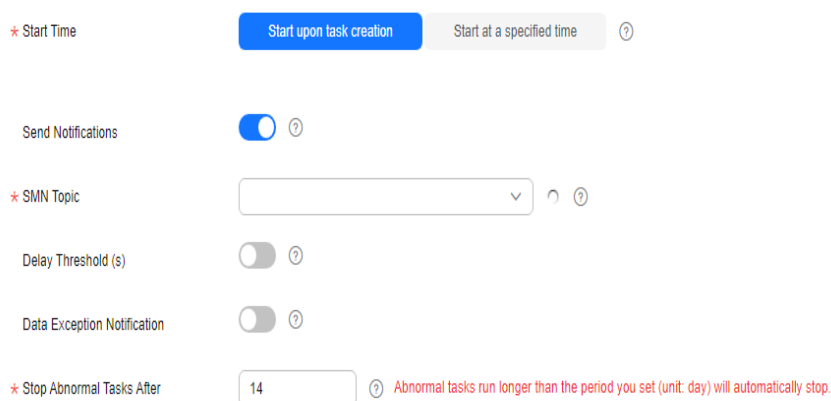



Table 4-247 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.23 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 4-248 Supported databases

Source DB	Destination DB
GaussDB primary/standby	Kafka 0.11 or later

Supported Synchronization Objects

Table 4-249 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-249 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenario: Incremental synchronization ● Supported fields: BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and TINYINT ● Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-250](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-250 Database user permission

Type	Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none">• The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables.• The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database.• Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-251 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If the DRS instance type is set to primary/standby, duplicate or inconsistent data may occur in tables without primary keys when a primary/standby switchover is performed. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: select working_version_num();

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During migration of table-level objects, you are not advised to rename the tables. • Replication of interval partition tables is not supported. • The name of a primary key column cannot be changed. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-203 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-252 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-204 Synchronization instance details

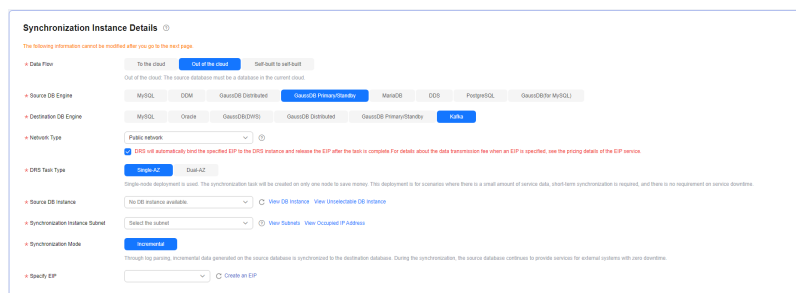


Table 4-253 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select Kafka .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization mode	<ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>

Parameter	Description
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 4-205 Task type




Table 4-254 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-206 AZ</p> <p>The screenshot shows two rows of radio buttons for AZ selection. The first row is labeled 'Primary AZ' and has buttons for 'az1' (selected), 'az2', 'az3', and 'az7'. The second row is labeled 'Standby AZ' and has buttons for 'az1' (selected), 'az2', 'az3', and 'az7'.</p>

- Enterprise Project and Tags

Figure 4-207 Enterprise Project and Tags




Table 4-255 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-208 Source database information

Source Database

DB Instance Name

Database Username

Database Password

Table 4-256 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-209 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 4-257 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-210 Synchronization mode

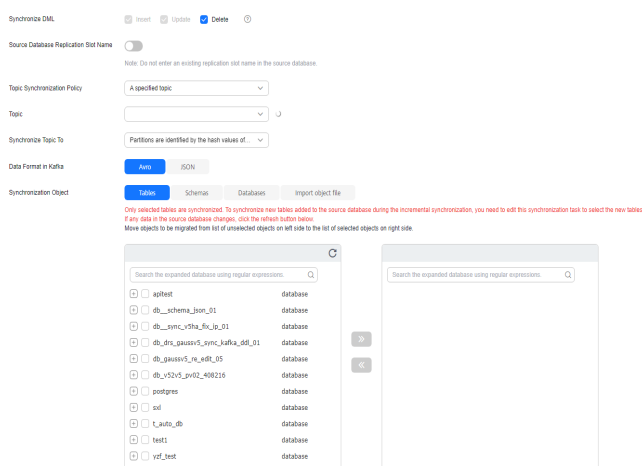



Table 4-258 Synchronization object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy is set to A specified topic .
Topic Name Format	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>

Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • Partitions are differentiated by the hash values of database_name.schema_name.table_name: This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved. • Partitions are identified by the hash values of the primary key: This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.schema_name.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Data Format in Kafka	<p>Select the format of data delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro: A binary encoded format that is efficient. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-211 Task startup settings

Table 4-259 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.24 From GaussDB Primary/Standby to GaussDB Distributed

Supported Source and Destination Databases

Table 4-260 Supported databases

Source DB	Destination DB
GaussDB primary/standby	GaussDB distributed NOTE The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

Table 4-261 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-261 Supported synchronization objects

Type	Restrictions
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the structure, sequence, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Restrictions
	<ul style="list-style-type: none">- The source table cannot contain a UNIQUE constraint. Otherwise, an error will be reported when data is written to the table structure of GaussDB Distributed, resulting in a task failure.- The database name, schema name, and table name cannot contain special characters /<.>\\` ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-262](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-262 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If <code>gs_loader</code> is used to create system catalogs (such as <code>public.pgxc_copy_error_log</code> and <code>public.gs_copy_summary</code>) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-263 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the

Type	Constraints
	<p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance.

Type	Constraints
	<ul style="list-style-type: none"> - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not compare data during full synchronization. After full synchronization is complete, the comparison is automatically triggered. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-212 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 4-264 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-213 Synchronization instance details

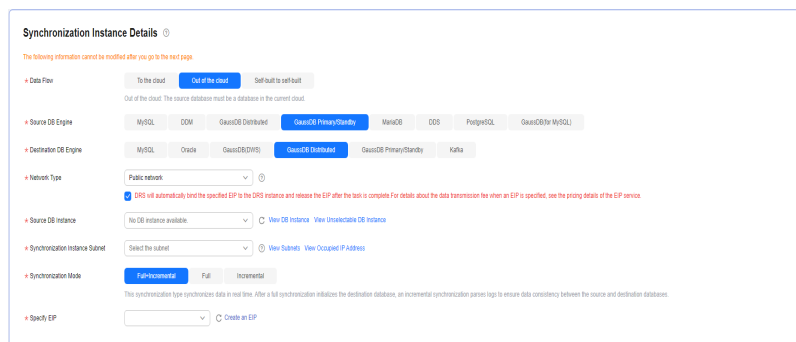


Table 4-265 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB Distributed .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB primary/standby instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-214 Task type




Table 4-266 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-215 Enterprise Project and Tags




Table 4-267 Enterprise Project and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-216 Source database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 4-268 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-217 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 4-269 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-218 Synchronization mode

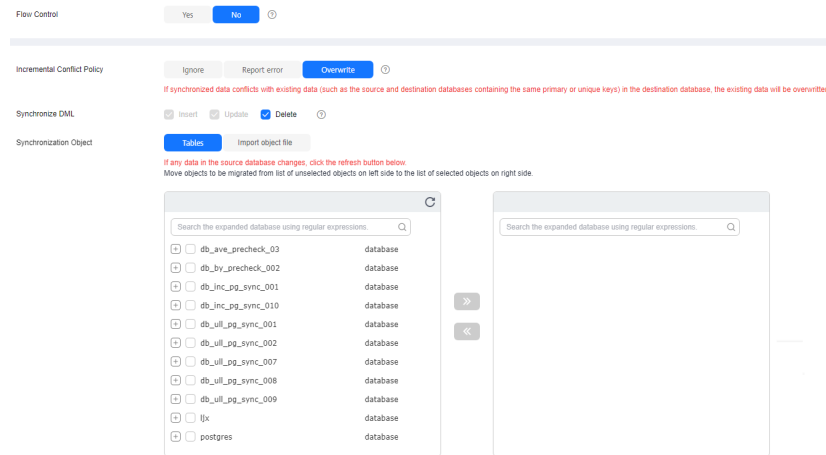
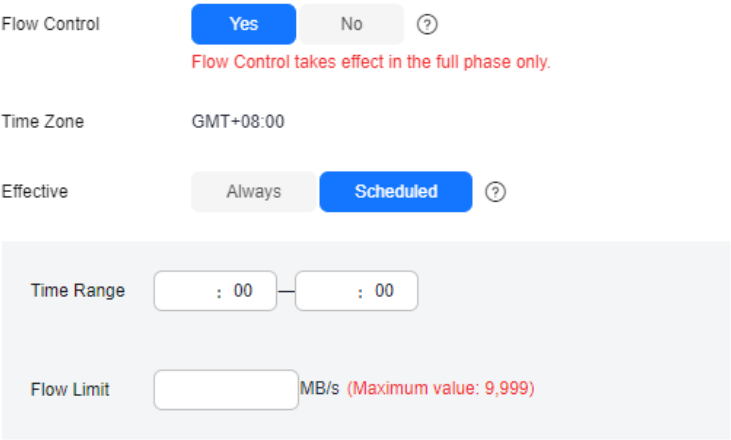



Table 4-270 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-219 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process.● Overwrite Conflicting data will be overwritten.● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-220 Task startup settings

Table 4-271 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.25 From GaussDB Primary/Standby to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-272 Supported databases

Source DB	Destination DB
GaussDB primary/standby	GaussDB primary/standby NOTE The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

[Table 4-273](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-273 Supported synchronization objects

Type	Restrictions
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the structure, sequence, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Restrictions
	<ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-274](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-274 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		<p>Connection and Port Description for Incremental Synchronization from GaussDB.</p>	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-275 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Select full or full+incremental synchronization. If a table structure already exists in the destination database, ensure that the table structure is the same as that on the source database. If column processing is performed, ensure that the table structure is the same as that after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the

Type	Constraints
	<p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not compare data during full synchronization. After full synchronization is complete, the comparison is automatically triggered. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-221 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 4-276 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-222 Synchronization instance details

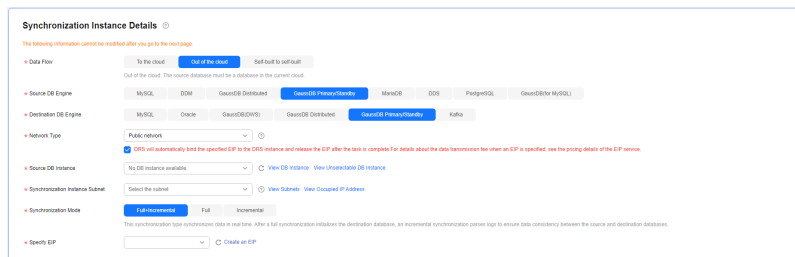


Table 4-277 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB primary/standby instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-223 Task type




Table 4-278 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-224 Enterprise Project and Tags




Table 4-279 Enterprise Project and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-225 Source database information

Source Database

DB Instance Name

Database Username


Database Password 

Table 4-280 Source database settings

Parameter	Description
DB Instance Name	The GaussDB primary/standby instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-226 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 4-281 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-227 Synchronization mode

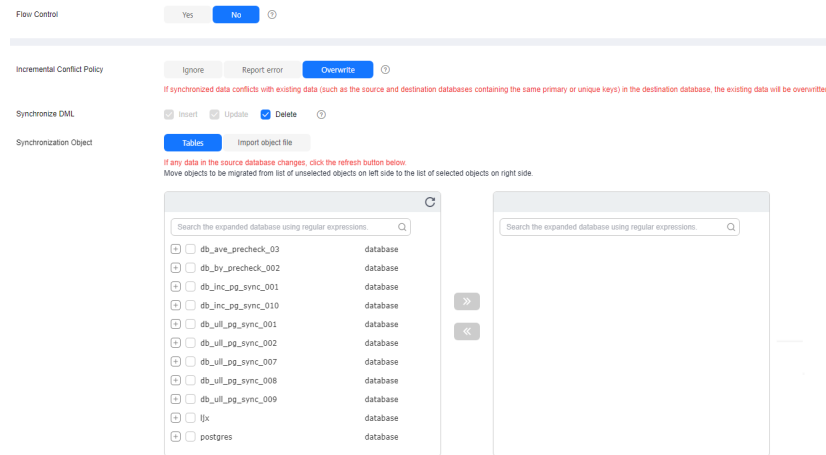
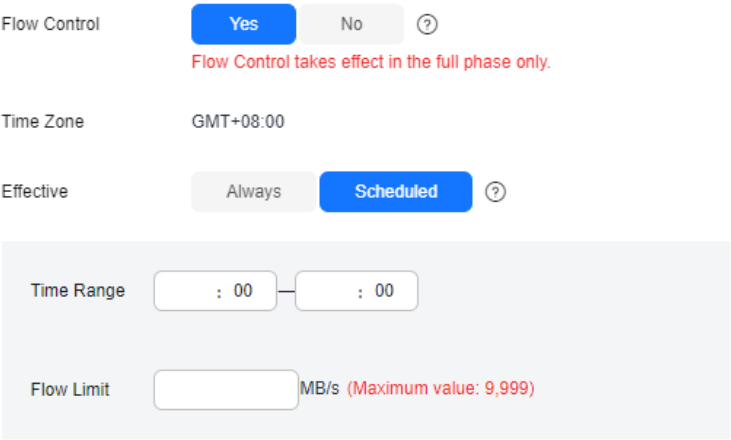



Table 4-282 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-228 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process.● Overwrite Conflicting data will be overwritten.● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-229 Task startup settings

Table 4-283 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.26 From GaussDB Primary/Standby to PostgreSQL

Supported Source and Destination Databases

Table 4-284 Supported databases

Source DB	Destination DB
GaussDB Primary/Standby	<ul style="list-style-type: none">On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-285 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-285 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● The following field types are not supported: SMALLDATETIME, REETIME, ABSTIME, TID, XID, CID and OID ● Table-level synchronization is supported. <ul style="list-style-type: none"> - Table data and sequence values can be synchronized. - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<>\\` ! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 4-286](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-286 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	The user must have the following minimum permissions: <ul style="list-style-type: none"> ● Database permission: CONNECT ● Schema permission: USAGE ● Table permission: INSERT, UPDATE, DELETE, and SELECT ● Sequence permission: UPDATE 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-287 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The destination database, schemas, and table object structures must be created in advance. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Constraints
	<ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● When GaussDB is not compatible with PostgreSQL, the processing precision of the date data type in the GaussDB database may be different from that in the PostgreSQL database. As a result, data of the two databases is inconsistent.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-230 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 4-288 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-231 Synchronization instance details

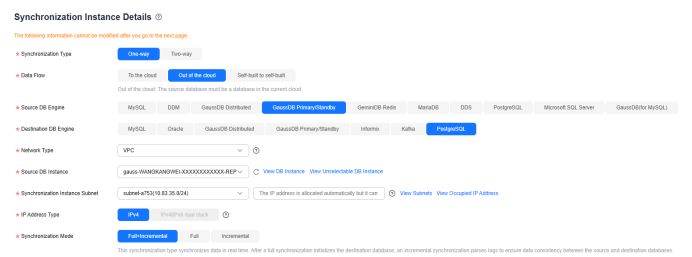


Table 4-289 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select PostgreSQL .
Network Type	<p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
Source DB Instance	The primary/standby GaussDB instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-232 AZ



Table 4-290 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-233 Enterprise Project and Tags




Table 4-291 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-234 Source database information

Source Database

DB Instance Name gauss-WANGKANGWEI-XXXXXXXXXXXX-REPL
(10.83.35.30:8000,10.83.35.107:8000,10.83.35.2:8000)

Database Username

Database Password

Test Connection
This button is available only after the replication instance is created successfully.

Table 4-292 Source database settings

Parameter	Description
DB Instance Name	The primary/standby GaussDB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 4-235 Destination database information

Destination Database

Select Connection

VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the destination database, and related parameters have been correctly configured.

This button is available only after the replication instance is created successfully.

Table 4-293 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 4-236 Synchronization mode

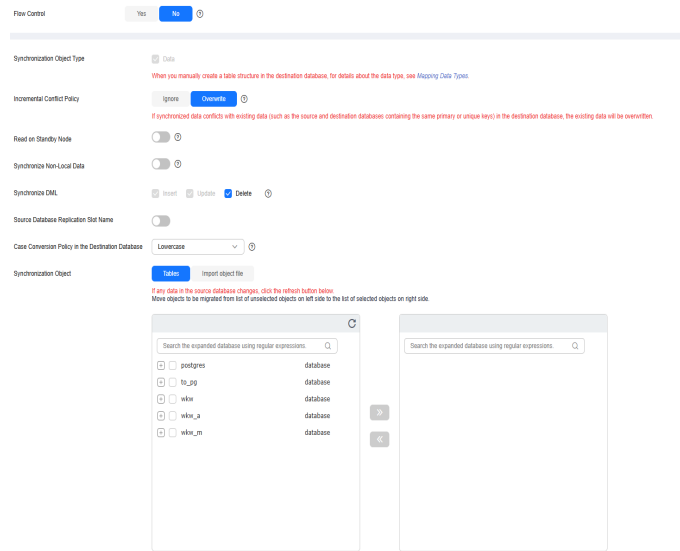
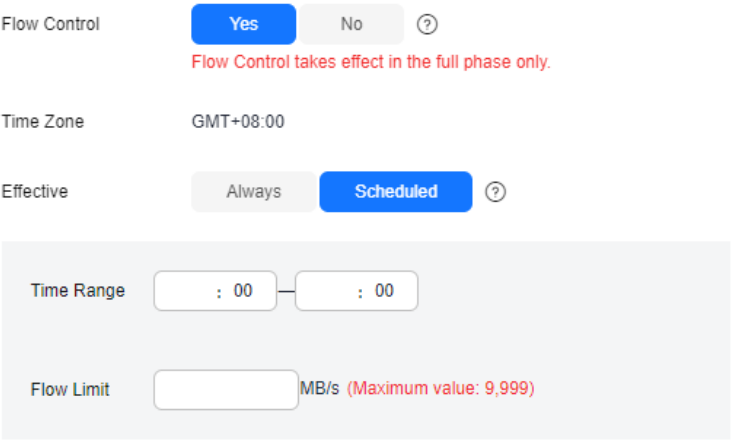



Table 4-294 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-237 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	Type of objects for full synchronization. Data is mandatory.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.
Case Conversion Policy in the Destination Database	Case of schema names, table names, and column names can be converted. If you have specified the mapping name when selecting objects to be synchronized, ignore this parameter.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-238 Task startup settings

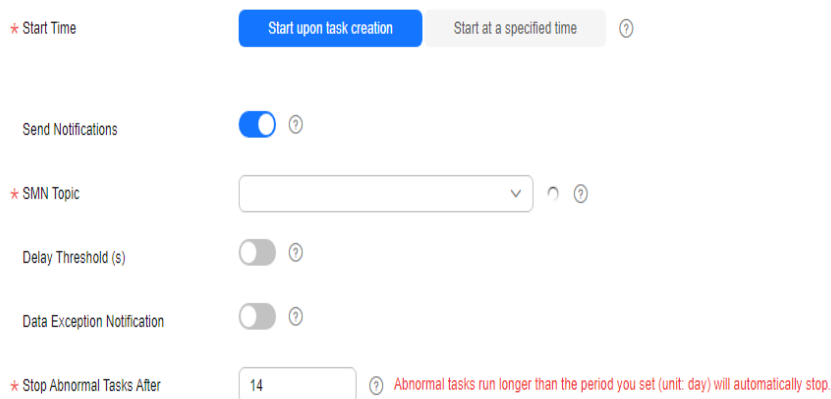



Table 4-295 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.27 From GaussDB(for MySQL) to MySQL

Supported Source and Destination Databases

Table 4-296 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> GaussDB(for MySQL) Primary/Standby 	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-297 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none">• The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT.• The destination database user must have the SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES permissions. By default, the root account of the RDS for MySQL instance has the preceding permissions. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, primary key indexes, unique indexes, common indexes, store procedures, views, and functions can be synchronized. ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● During database name mapping, function objects, views, and stored procedures may fail to be created because they are referenced by a database table name before the mapping. DRS ignores the error reported during the creation of these objects in the full synchronization phase. As a result, these objects are lost in the destination database. In the incremental synchronization phase, if database name mapping is required, DDL operations (including CREATE, MODIFY, and DELETE) related to function objects, views, and stored procedures are not synchronized to the destination database. ● If the database table name contains characters other than letters, digits, and underscores (<code>_</code>), or the mapped database table name contains hyphens (<code>-</code>) and number signs (<code>#</code>), the name length cannot exceed 42 characters. ● Only MyISAM and InnoDB tables can be synchronized. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the MySQL source database server_id must be set. The value ranges from 1 to 4294967296. ● The names of the source databases, tables, and views cannot contain non-ASCII characters, or the following characters: '<>/' ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized to RDS. ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● Database mapping does not support views or stored procedures. If the source database contains views or stored procedures, the synchronization may fail.
Destination database	<ul style="list-style-type: none"> ● The destination database is running properly. ● The destination database must have sufficient disk space. ● If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. ● The character set of the destination database must be the same as that of the source database. ● The time zone of the destination database must be the same as that of the source database. ● If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss. ● The source and destination databases cannot contain tables that have the same names but do not have primary keys. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The source database does not support point-in-time recovery (PITR). ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. ● DDL operations are not supported during full synchronization. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX,

Type	Restrictions
	<p>ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required.</p> <ul style="list-style-type: none"> - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-239 Synchronization task information

▲ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-298 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-240 Synchronization instance details

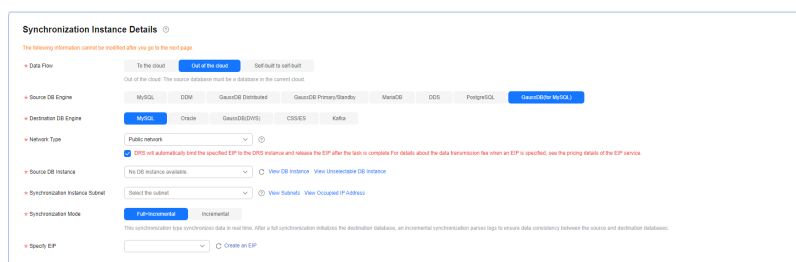


Table 4-299 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select MySQL .

Parameter	Description
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The GaussDB(for MySQL) instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Incremental Through log parsing, incremental data generated on the source is synchronized to the destination in real time.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task type

Figure 4-241 Task type




Table 4-300 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-242 Enterprise Project and Tags




Table 4-301 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-243 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-302 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the system and will be permanently deleted after the task is deleted.

Figure 4-244 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 4-303 Destination database settings

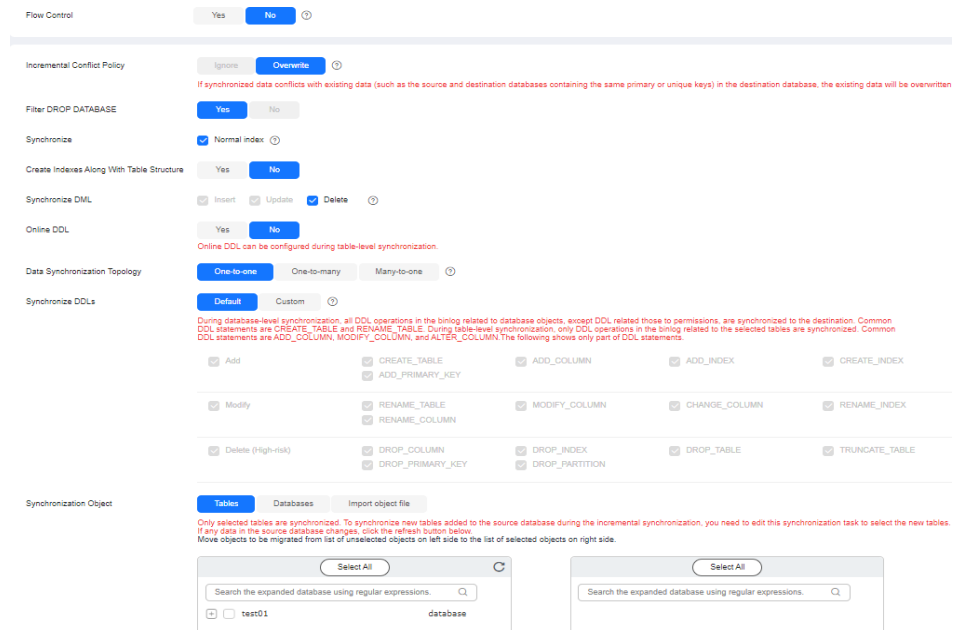
Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.

 NOTE

The IP address, username, and password of the destination database are encrypted and stored in the system. After the task is deleted, the information is permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization object. Click **Next**.

Figure 4-245 Synchronization mode



Flow Control: Yes **No**

Incremental Conflict Policy: Ignore **Overwrite**
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Filter DROP DATABASE: Yes **No**

Synchronize: Normal index

Create Indexes Along With Table Structure: Yes **No**

Synchronize DML: Insert Update Delete

Online DDL: Yes **No**
Online DDL can be configured during table-level synchronization.

Data Synchronization Topology: **One-to-one** One-to-many Many-to-one

Synchronize DDLs: **Default** Custom
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. The following shows only part of DDL statements.

<input checked="" type="checkbox"/> Add	<input checked="" type="checkbox"/> CREATE_TABLE <input checked="" type="checkbox"/> ADD_PRIMARY_KEY	<input checked="" type="checkbox"/> ADD_COLUMN	<input checked="" type="checkbox"/> ADD_INDEX	<input checked="" type="checkbox"/> CREATE_INDEX
<input checked="" type="checkbox"/> Modify	<input checked="" type="checkbox"/> RENAME_TABLE <input checked="" type="checkbox"/> RENAME_COLUMN	<input checked="" type="checkbox"/> MODIFY_COLUMN	<input checked="" type="checkbox"/> CHANGE_COLUMN	<input checked="" type="checkbox"/> RENAME_INDEX
<input checked="" type="checkbox"/> Delete (high-risk)	<input checked="" type="checkbox"/> DROP_COLUMN <input checked="" type="checkbox"/> DROP_PRIMARY_KEY	<input checked="" type="checkbox"/> DROP_INDEX <input checked="" type="checkbox"/> DROP_PARTITION	<input checked="" type="checkbox"/> DROP_TABLE	<input checked="" type="checkbox"/> TRUNCATE_TABLE

Synchronization Object: **Tables** Databases Import object file
Only selected tables are synchronized. To synchronize new tables added to the source database during the incremental synchronization, you need to edit this synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

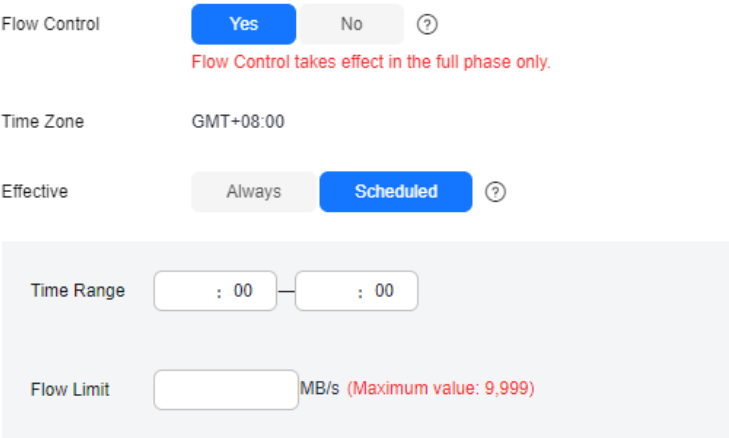
Search the expanded database using regular expressions.

test01 database

Select All

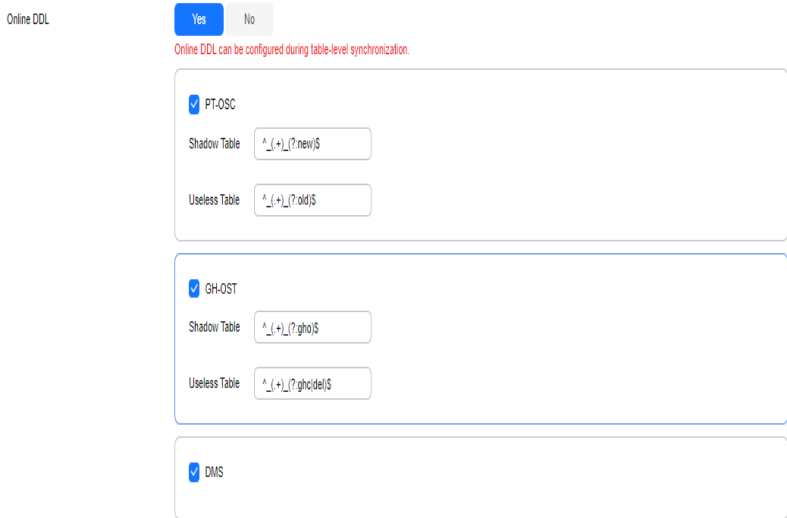
Search the expanded database using regular expressions.


Table 4-304 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-246 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database.● Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none">● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization.● If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	Select whether to synchronize indexes based on the service requirements.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none">● Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase.● No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Data Synchronization Topology	Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies .
Synchronize DDLs	<p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OCS, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 4-247 Online DDL</p>  <ul style="list-style-type: none"> No: Table-level synchronization does not support Online DDL synchronization.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If you need to set data filtering, click **Data Filtering** and set related filtering rules.
- If you need to add additional columns, click the **Additional Columns** tab, click **Add** in the **Operation** column, and enter the column name and the operation type.

For details about related operations, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-248 Task startup settings

Table 4-305 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.28 From GaussDB(for MySQL) to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-306 Supported databases

Source DB	Destination DB
Primary/standby GaussDB(for MySQL) instances	GaussDB(DWS) cluster

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-307 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. • The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, common indexes (B-Tree indexes), and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, and unique indexes cannot be synchronized. ● Comment is supported in full synchronization mode, but not supported in incremental synchronization mode. ● The following data types are not supported: XML, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon. ● Only MyISAM and InnoDB tables can be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● Replication tables without primary keys of GaussDB(DWS) cannot be synchronized. If a table to be synchronized is a replication table without a primary key in GaussDB(DWS), the task will fail. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● The source database server-id must be set to a value ranging from 2 to 4294967296. ● The source database name and table name cannot contain non-ASCII characters, or special characters '<>/\
Destination database	<ul style="list-style-type: none"> ● The destination database is running properly. ● The destination database must have sufficient disk space. ● The time zone of the destination database must be the same as that of the source database.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common association: Indexes reference tables. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. ● If there is a unique key when a table without a primary key is synchronized, data may fail to be written due to data conflicts. ● If the character sets of the source and destination databases are different, data may be inconsistent or synchronization may fail. ● If the data types are incompatible, the synchronization may fail. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails.

Type	Restrictions
	<ul style="list-style-type: none"> • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • During real-time synchronization, you can change the port number. If the synchronization task fails after the port is changed, you can retry the synchronization task. • During real-time synchronization, the IP address, account, and password cannot be changed. • During incremental synchronization, the following DDL operations are supported by default: <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, ALTER_COLUMN, DROP_CONSTRAINT, ADD_CONSTRAINT, CREATE_INDEX, DROP_INDEX, RENAME_INDEX, and RENAME_COLUMN. You can select the DDL operations to be synchronized on the object selection page as required. - If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. - In many-to-one scenarios, you are advised to synchronize ADD_COLUMN only. Other DDL synchronization may cause task failures or data inconsistency due to destination table changes. - In many-to-one scenarios, when running ADD_COLUMN, ensure that the types of columns added to each table are the same. Otherwise, the task may fail. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - When an index is creating using a DDL statement, if the table name and index name in the statement contain more than 63 characters, duplicate names may occur. As a result, the index fails to be created. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. - When a DDL operation is performed in the incremental phase, if the destination table is not found, the DDL operation will be ignored. - In the incremental phase, if CHANGE COLUMN is performed in the source database to modify a column and the column is a distribution column in destination

Type	Restrictions
	<p>GaussDB(DWS), the statement will be ignored because GaussDB(DWS) does not support distribution column modification.</p> <ul style="list-style-type: none"> - In the incremental phase, the RENAME INDEX operation is not supported because the index rules of the source and destination databases are different. - In the incremental phase, you are not advised to combine CHAR(0) with other characters for data synchronization. For example, inserting CONCAT('a',CHAR(0),'b') may cause data inconsistencies. - In the incremental phase, you are not advised to combine CHAR(34) (double quotation marks) and CHAR(92) (backslash) with other characters for JSON data synchronization. For example, inserting JSON_OBJECT('\$\$.0', CONCAT('a', CHAR(34), 'b')) may cause data inconsistencies because an extra backslash (\) will be added after it is synchronized to the destination end. <ul style="list-style-type: none"> • During an incremental synchronization, database-level synchronization does not support online DDL, and table-level synchronization supports only online DDL generated by Alibaba Cloud DMS. • During an incremental synchronization, tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-249 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name:

Description:

0/256

Table 4-308 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-250 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud Out of the cloud Self-built to self-built

Source DB Engine: MySQL EDM GaussDB Distributed GaussDB Primary/Standby MySQL DDS PostgreSQL GaussDB for MySQL

Destination DB Engine: MySQL Oracle GaussDB(for MySQL) CSSES Kafka

Network Type: Public network Private network

Source DB Instance: No DB instance available New DB Instance Unavailable DB Instance

Synchronization Instance Subnet: Select the subnet New Subnets View Occupied IP Address

Synchronization Mode: Full Synchronization

Specify EP: Create an EP

Table 4-309 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .

Parameter	Description
Destination DB Engine	Select GaussDB(DWS) .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	An available GaussDB(for MySQL) instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-251 Task type




Table 4-310 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-252 Enterprise Project and Tags




Table 4-311 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-253 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-312 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in the system and will be permanently deleted after the task is deleted.

Figure 4-254 Destination database information

Destination Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

Test Connection This button is available only after the replication instance is created successfully.

Table 4-313 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The IP address, username, and password of the destination database are encrypted and stored in the system. After the task is deleted, the information is permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization object. Click **Next**.

NOTE

Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Figure 4-255 Synchronization mode

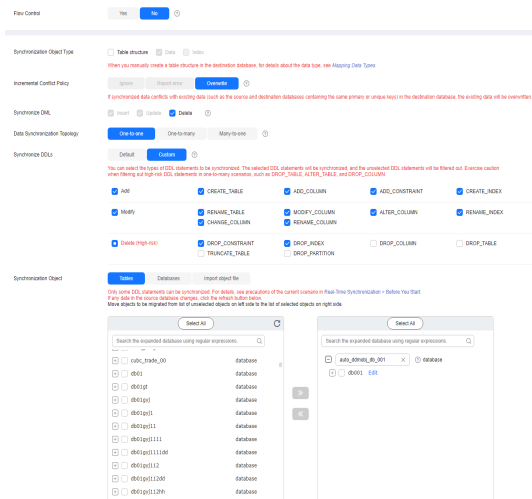
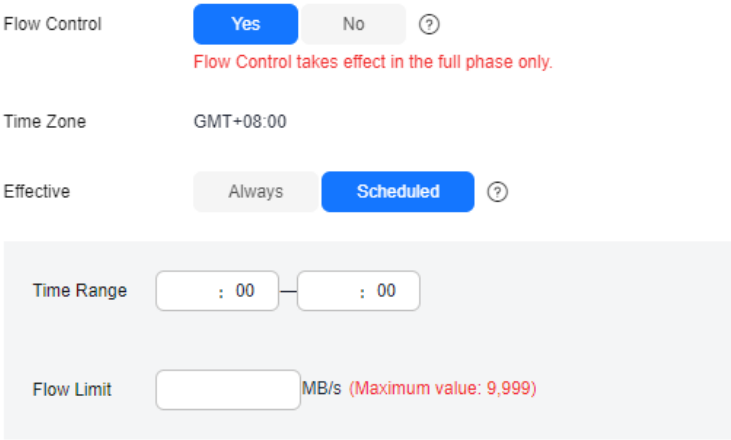



Table 4-314 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-256 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Data Synchronization Topology	<p>Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>

Parameter	Description
Synchronize DDLs	<p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Data Processing** page, select the table object to be processed, click **Add** in the **Operation** column, enter the column name, type, and operation type, confirm the information, and click **Next**. You can configure related rules by referring to [Processing Data](#).

Figure 4-257 Processing data

Additional Columns

You can use additional columns to avoid data conflicts in many-to-one scenarios.

Source Database Table	New Name	Column Name	Operation Type	Type	Operation
ts_server_atl@changep_002@tblc_0_1	ts_server_atl@changep_002@tblc_0_1	--	--	--	Add
ts_server_atl@changep_002@tblc_0_2	ts_server_atl@changep_002@tblc_0_2	--	--	--	Add
ts_server_atl@changep_002@tblc_0_3	ts_server_atl@changep_002@tblc_0_3	--	--	--	Add
ts_server_atl@changep_002@tblc_0_4	ts_server_atl@changep_002@tblc_0_4	--	--	--	Add

 **CAUTION**

- If you fill in a new column in *serverName@database@table* format, the new column and the primary key of the source table form a composite primary key.

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-258 Task startup settings

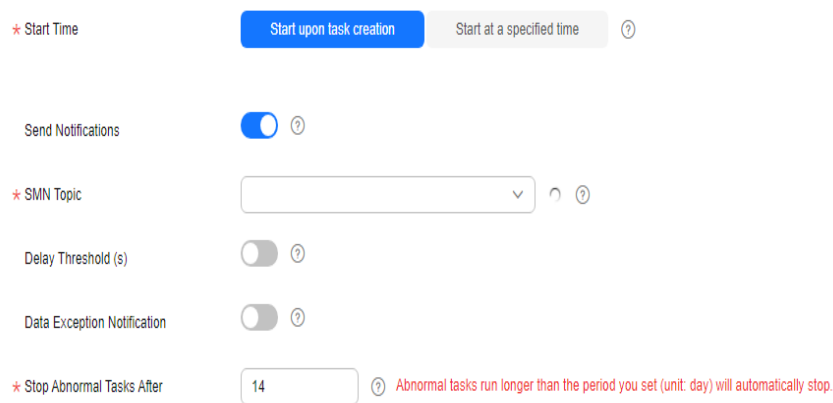



Table 4-315 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.29 From GaussDB(for MySQL) to Kafka

Supported Source and Destination Databases

Table 4-316 Supported databases

Source DB	Destination DB
Primary/standby GaussDB(for MySQL) instances	Kafka

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-317 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT.
Synchronization object	<ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>
Source database	<ul style="list-style-type: none"> • The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. • During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. • If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. • If the expire_logs_days value of the source database is set to 0, the synchronization may fail. • Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. • The server_id value of the source MySQL database must be in the range from 1 to 4294967296. • The database and table names in the source database cannot contain non-ASCII characters, or special characters '<>\'/\'

Type	Restrictions
Destination database	<ul style="list-style-type: none"> • The destination database is a Kafka database. • You are advised to set auto.create.topics.enable of Kafka to false.
Precautions	<ul style="list-style-type: none"> • If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. • If the data types are incompatible, the synchronization may fail. • If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. • Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. • The source database does not support point-in-time recovery (PITR). • Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During synchronization of table-level objects, renaming tables is not recommended. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-259 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] .. v
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu] .. v

* Task Name: DRS-5678 ⓘ

Description: [Text area] ⓘ
0/256

Table 4-318 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-260 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud Self-host to self-host
Out of the cloud: The source database must be a database in the current cloud.

• Source DB Engine MySQL, ODB, GaussDB Distributed, GaussDB Primary/Standby, MaxCompute, DDS, PostgreSQL [Search for MySQL](#)

• Destination DB Engine MySQL, Oracle, GaussDB(DWS), CSSS3, Auto

• Network Type Public network ⓘ
 DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

• DRS Task Type Single AZ Dual AZ
Single AZ: Single AZ deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, load data synchronization is required, and there is no requirement on service downtime.

• Source DB Instance No DB instance available ⓘ [View DB Instance](#) [View Unavailable DB Instance](#)

• Synchronization Instance Subnet Select the subnet ⓘ [View Subnet](#)

• Synchronization Mode Full-Refresh Incremental
This synchronization type synchronizes data in real time, after a full synchronization initializes the destination database, an incremental synchronization process logs to ensure data consistency between the source and destination databases.

• Specify EP ⓘ Create an EP

Table 4-319 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select Kafka .
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The GaussDB(for MySQL) instance you created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<p>Available options: Full+Incremental and Incremental</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 4-261 Task type





Table 4-320 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-262 AZ</p> 

- Enterprise Project and Tags

Figure 4-263 Enterprise Project and Tags




Table 4-321 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-264 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-322 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-265 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

✔ Test successful

Table 4-323 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-266 Synchronization mode

Flow Control Yes No ⓘ

Synchronize DML: Insert Update Delete ⓘ

Topic Synchronization Policy:

Topic: C

Synchronize Topic To:

Data Format in Kafka:

Synchronization Object:

Only selected tables are synchronized. To synchronize new tables added to the source database during the incremental synchronization, you need to edit this synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. More objects to be migrated from list of expected objects on left side to the list of selected objects on right side.

Select All

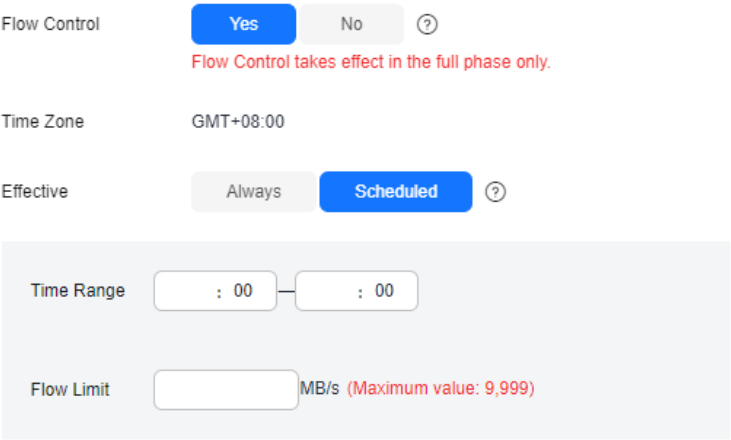
Search the expanded database using regular expressions

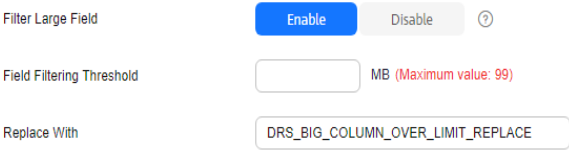
- db_1_zpg_multidatabase_001 database
- db_1_xq234ka_dtable_at_need0_ database
- db_2_zpg_multidatabase_001 database
- db_mingfang_nftx_002 database
- db_news_check_001 database
- db_xms_and_xms_001 database
- db_bjeditreport_001 database
- db_bjg20230913030918 database
- db_dhndashl_gym_004 database
- db_eft_dtable_003 database
- db_eft_dtable_004 database
- db_eft_dtable_004 database

Select All


Search the expanded database using regular expressions

Table 4-324 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-267 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Filter Large Field	<p>Indicates whether to use large field filtering to process special fields (blob, mediumblob, longblob, varbinary, mediumtext and longtext) in a synchronization table.</p> <ul style="list-style-type: none"> Enable You need to set Field Filtering Threshold and Replace With. If the size of a field exceeds the threshold, the value is replaced based on a specified character. Note that large field filtering is used to replace the value of a field, not the entire DML record. If a DML record contains many large fields, the size of only some of these fields exceeds the filtering threshold, and the accumulated value of other fields that do not exceed the filtering threshold is greater than the value of request.max.size, when data is written to Kafka, the size of the message body in the destination Kafka database may still exceed the upper limit, resulting in a DRS error. Figure 4-268 Setting large field filtering  Disable Large fields are not filtered.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>
Topic Synchronization Policy	<p>Topic synchronization policy. You can select A specific topic or Auto-generated topics.</p>

Parameter	Description
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal. If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.
Number of Partitions	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.
Replication Factor	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.
Synchronize Topic To	The policy for synchronizing topics to the Kafka partitions. <ul style="list-style-type: none"> ● If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. ● If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic.
Data Format in Kafka	Select the data format to be delivered from GaussDB(for MySQL) to Kafka. <ul style="list-style-type: none"> ● Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. ● JSON: JSON message format, which is easy to interpret but takes up more space. ● JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. For details, see Kafka Message Format .

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-269 Task startup settings





Table 4-325 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.30 From GaussDB(for MySQL) to CSS/ES

Supported Source and Destination Databases

Table 4-326 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> Primary/standby GaussDB(for MySQL) instances 	<ul style="list-style-type: none"> ElasticSearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9 and 7.10

 **NOTE**

Only whitelisted users can use this function.

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
 - Data-Level Comparison
To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-327 Precautions

Type	Restrictions
Database permissions	<p>Minimum permission requirements for full plus incremental synchronization:</p> <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the following permissions: READ and WRITE
Synchronization object	<ul style="list-style-type: none"> • The table data can be synchronized. • Databases, views, indexes, constraints, functions, stored procedures, triggers, and events cannot be synchronized. • The system database and event status cannot be synchronized. • Tables whose primary keys are of the FLOAT type cannot be synchronized. • Tables that do not have primary keys cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>/' ● The column name in the source database cannot contain non-ASCII characters or the following characters: '.' ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● During an incremental synchronization, the server_id value of the source database must be specified. The value of server_id ranges from 1 to 4294967296. ● Enable skip-name-resolve for the source database to reduce the possibility of connection timeout. ● Enable the Global Transaction Identifier (GTID) of the source database. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. ● The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. ● The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail.
Destination database	<ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored to a point in time when a full synchronization was being performed. ● Convert the value range of the source database time field that is not supported by the destination database to null. ● The strings padded with zeros in the source database may be truncated because the source database uses the fixed-length binary data type, and the destination database uses the variable-length data type. ● The binary value is encrypted using Base64 and then written to the destination database. ● If no time zone is specified for the source database, specify the time zone for synchronizing the datetime type to the destination database. ● All table field names are converted to lowercase letters. ● If the <code>_id</code> field of the destination database is generated using multiple columns in the source database, separate these columns with colons (:). ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● During the synchronization, do not create a database named ib_logfile in the source.

Type	Restrictions
	<ul style="list-style-type: none"> • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable upload is supported, but data may be repeatedly inserted into non-transactional tables that do not have primary keys when the server system breaks down. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-270 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-328 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-271 Synchronization instance details

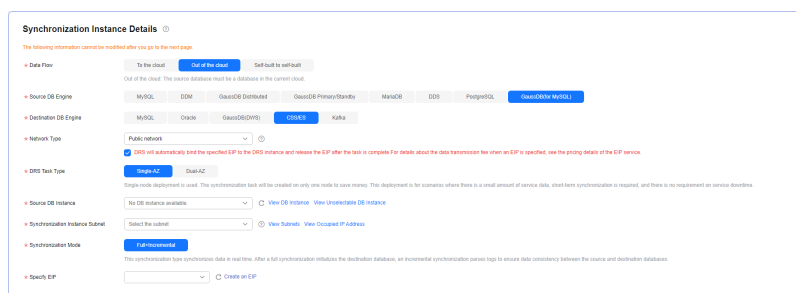


Table 4-329 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select CSS/ES .

Parameter	Description
Network Type	<p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Source DB Instance	The GaussDB(for MySQL) instance you created.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	– Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 4-272 Task type





Table 4-330 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-273 AZ</p> 

- Enterprise Project and Tags

Figure 4-274 Enterprise Project and Tags




Table 4-331 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-275 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-332 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-276 Destination database information

Destination Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-333 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. Only .cer and .pem certificates are supported. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-277 Synchronization mode

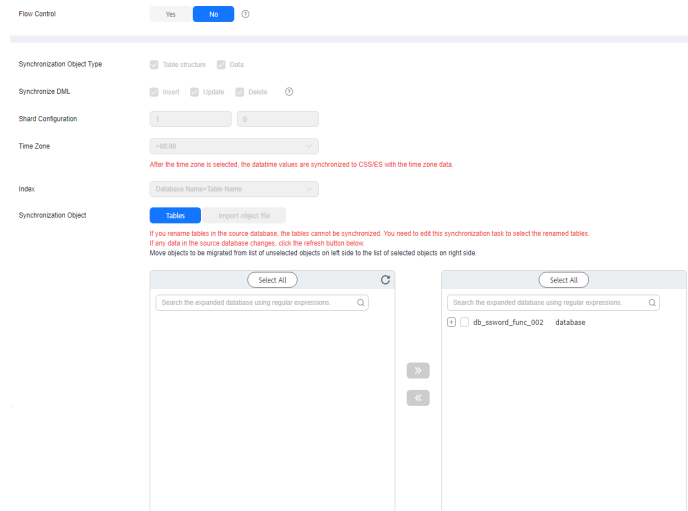
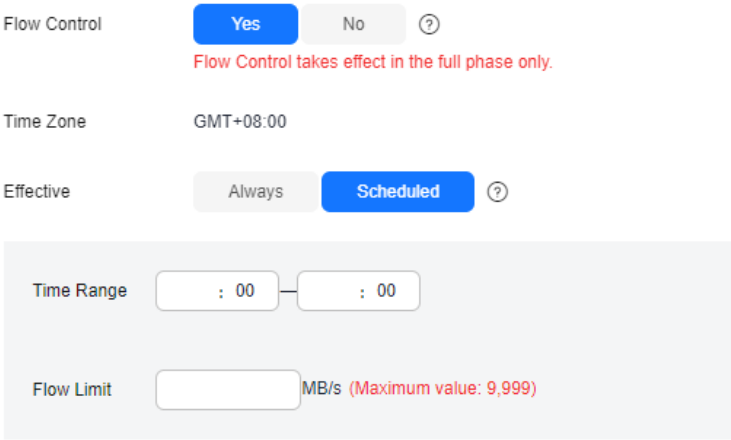



Table 4-334 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-278 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Shard Configuration	<p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p>
Time Zone	<p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p>
Index Name	<ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects. For details, see Editing a Synchronization Task. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

 **NOTE**

The column processing of the synchronization task from GaussDB(for MySQL) to CSS/ES supports only column filtering and does not support column mapping.

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-279 Task startup settings

Table 4-335 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.31 From GaussDB(for MySQL) to Oracle

Supported Source and Destination Databases

Table 4-336 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> Primary/standby GaussDB(for MySQL) instances 	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-337 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT The destination database user must have the following permissions: ALTER ANY INDEX, ALTER ANY TABLE, ALTER SESSION, ANALYZE ANY, COMMENT ANY TABLE, CREATE ANY INDEX, CREATE ANY TABLE, CREATE SESSION, DELETE ANY TABLE, DROP ANY TABLE, INSERT ANY TABLE, SELECT ANY TABLE, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, UPDATE ANY TABLE, and RESOURCE roles.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Full synchronization supports the synchronization of data, table structures, and indexes. ● Incremental synchronization supports only data synchronization. ● Geography data types such as geometry, geometrycollection, linestring, multilinestring, multipoint, point and polygon are not supported. ● Views, constraints, functions, stored procedures, triggers, and events cannot be synchronized. ● The system database and event status cannot be synchronized. ● The destination Oracle database does not support empty strings, so the object to be synchronized cannot contain empty strings. ● The maximum number of columns supported by the source GaussDB(for MySQL) database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>/' ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● During an incremental synchronization, the source database server_id must be set to a value ranging from 1 to 4294967296. ● Enable skip-name-resolve for the source database to reduce the possibility of connection timeout. ● Enable the Global Transaction Identifier (GTID) of the source database. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. ● The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. ● The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail. ● Do not set lower_case_table_names to 0.
Destination database	<ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the default value of the time field in the source database is all 0s, it will be converted to 1970-01-01 00:00:00. ● If the precision of source database decimal data type exceeds 38, the data will be truncated because the maximum precision of the destination database number data type is 38. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● The source database does not support fields whose column type is binary and the length is 0 (that is, binary(0)). This is because the binary type of GaussDB(for MySQL) is mapped to the raw type when being synchronized to Oracle using DRS. However, in Oracle, the length of the raw type cannot be set to 0. ● All table field names are converted to uppercase letters. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.

Type	Restrictions
	<ul style="list-style-type: none"> • During the synchronization, do not create a database named ib_logfile in the source. • During an incremental synchronization, do not perform the point-in-time recovery (PITR) operation on the source database. • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable upload is supported, but data may be repeatedly inserted into non-transactional tables that do not have primary keys when the server system breaks down. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-280 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0256

Table 4-338 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-281 Synchronization instance details

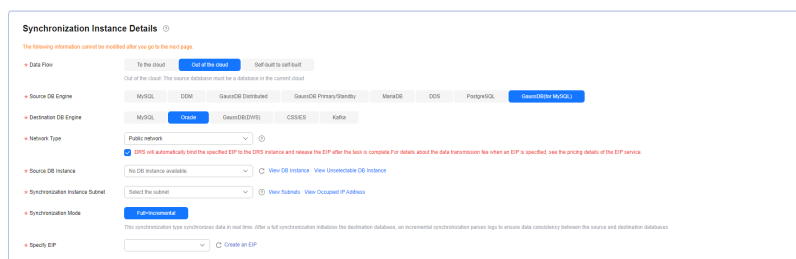


Table 4-339 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select Oracle .
Network Type	The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
Source DB Instance	The GaussDB(for MySQL) instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	– Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-282 Task type




Table 4-340 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-283 Enterprise Project and Tags




Table 4-341 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-284 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-342 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-285 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

Database Username

Database Password

SSL Connection

Table 4-343 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address to improve access performance.
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-286 Synchronization mode

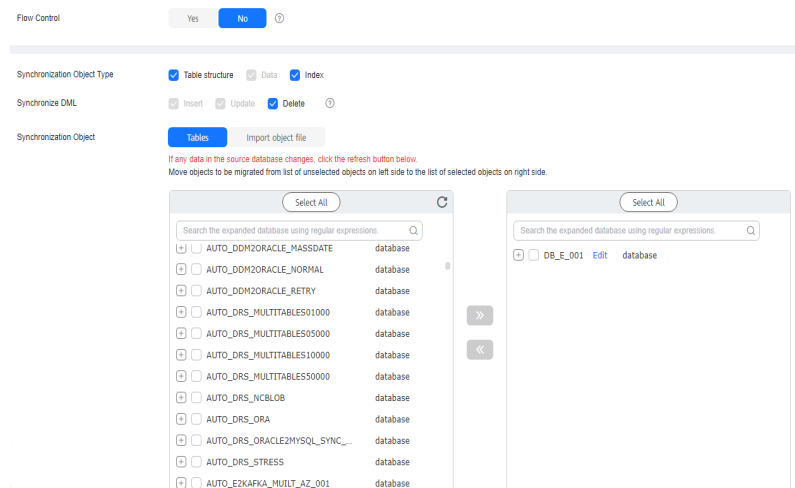
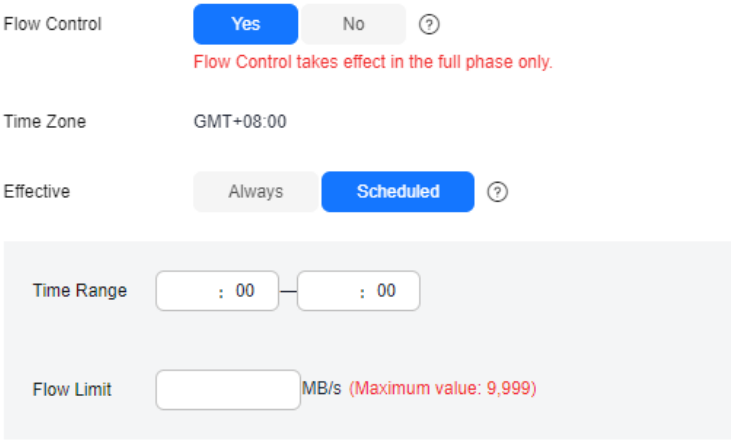



Table 4-344 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-287 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <p>For details about how to import an object file, see Importing Synchronization Objects.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-288 Task startup settings

The screenshot shows the following settings:


- Start Time:** Two radio buttons. The first, "Start upon task creation", is selected and highlighted in blue. The second, "Start at a specified time", is unselected.
- Send Notifications:** A toggle switch is turned on (blue).
- SMN Topic:** A dropdown menu is currently empty.
- Delay Threshold (s):** A toggle switch is turned off (grey).
- Data Exception Notification:** A toggle switch is turned off (grey).
- Stop Abnormal Tasks After:** An input field contains the number "14". A red note below it says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-345 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.32 From MariaDB to MariaDB

Supported Source and Destination Databases

Table 4-346 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">RDS for MariaDB 10.5	<ul style="list-style-type: none">On-premises MariaDB 10.5ECS-hosted MariaDB 10.5Other cloud MariaDB 10.5

Supported Synchronization Objects

Table 4-347 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-347 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. ● Only MyISAM and InnoDB tables can be synchronized. ● Events and triggers cannot be synchronized. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-348](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-348 Database account permission

Type	Full+Incremental
Source database user	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)
Destination database user	SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-349 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During the incremental synchronization, the binlog of the source MariaDB database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the expire_logs_days value of the source database is set to 0, the synchronization may fail. - Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. - During an incremental synchronization, the source MariaDB database server_id must be set to a value ranging from 1 to 4294967296. ● Source database object requirements: <ul style="list-style-type: none"> - The source database names cannot contain non-ASCII characters, or the following characters: <code>.'<>\'/\"</code> - The source table and view names cannot contain non-ASCII characters, or the following characters: <code>.'<>\'/\"</code> - The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● Destination database parameter requirements: <ul style="list-style-type: none"> - Data cannot be synchronized from a newer version database to an older version database. - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. - If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. - If the enforce_storage_engine parameter of the destination MariaDB database is set to InnoDB, DRS

Type	Restrictions
	<p>cannot synchronize the table structure and data whose storage engine is MyISAM to the destination MariaDB database. To synchronize table data whose storage engine is MyISAM, create a table structure on the destination database. (The storage engine can only be set to InnoDB due to the value restriction of the enforce_storage_engine parameter.)</p> <ul style="list-style-type: none"> ● Destination database object requirements: <ul style="list-style-type: none"> - The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal. - The destination DB instance must have sufficient storage space. - If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. - The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - If the sources and destinations are RDS instances, database mapping is required. - The source and destination databases cannot contain tables that have the same names but do not have primary keys. - If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. - The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. - If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. - If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection.

Type	Restrictions
	<ul style="list-style-type: none"> - Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. - The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. - The source database does not support point-in-time recovery (PITR). - The destination database cannot be restored to a point in time when a full synchronization was being performed. - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. - The partitioned table does not support column mapping. - Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - After a task is created, the destination database cannot be set to read-only. - If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement.

Type	Restrictions
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • Data inconsistency may occur when the MyISAM table is modified during a full synchronization. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • DDL statements can be synchronized during incremental synchronization. • Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • You can add additional objects during an incremental synchronization.
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison.

Procedure

This section describes how to use DRS to configure a MariaDB out-of-cloud synchronization task over a public network.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-289 Synchronization task information

Warning: Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page. The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

Task Name: [Text input: DRS-5678] ⓘ

Description: [Text area: 0/256] ⓘ

Table 4-350 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 4-290 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data File: [To the cloud] **Out of the cloud** [Self-built to self-built]

Out of the cloud: The source database must be a database in the current cloud.

• Source DB Engine: [MySQL] [ODM] [GaussDB Distributed] [GaussDB Primary/Standby] **GaussDB** [DDS] [PanguDB] [GaussDB for MySQL]

• Destination DB Engine: [MySQL]

• Network Type: [Public network] ⓘ

DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

• Source DB Instance: [No DB instance available] ⓘ [View DB Instance](#) [View Unavailable DB Instance](#)

• Synchronization Instance Status: [Select the status] ⓘ [View Subnets](#) [View Occupied IP Address](#)

• Synchronization Mode: [Full Synchronization] ⓘ

This synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization process logs to ensure data consistency between the source and destination databases.

• Security EP: [] ⓘ [Create an EP](#)

Table 4-351 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select MariaDB .
Destination DB Engine	Select MariaDB .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is the default value and is used as an example. <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	The RDS for MariaDB instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Full+Incremental : This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental , data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task type

Figure 4-291 Task type




Table 4-352 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-292 Enterprise Project and Tags




Table 4-353 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 4-293 Source database information

Source Database

DB Instance Name

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 4-354 Source database information

Parameter	Description
DB Instance Name	The RDS for MariaDB instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 4-294 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-355 Source database information

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-295 Synchronization objects

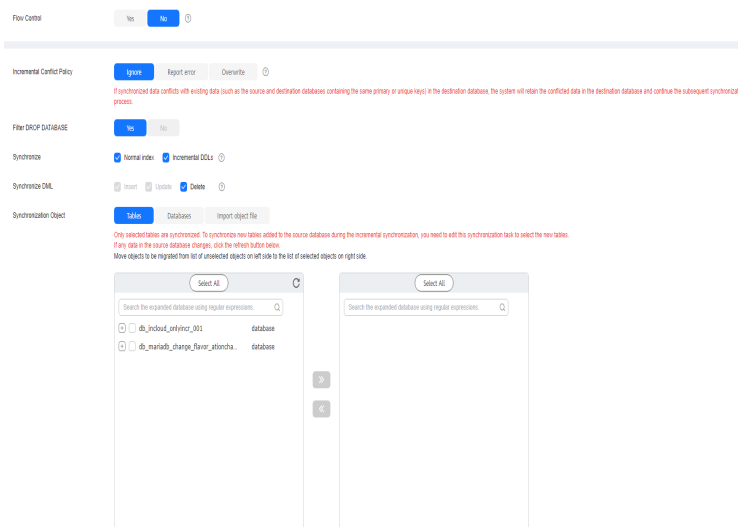
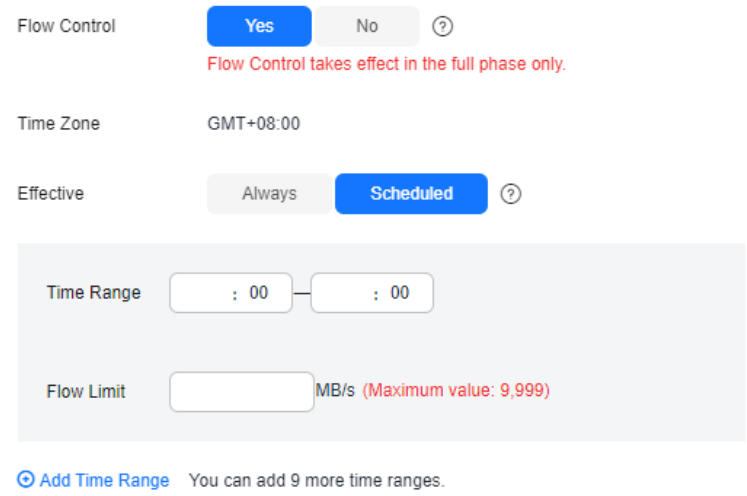



Table 4-356 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-296 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**. For details about how to configure related rules, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-297 Task startup settings





Table 4-357 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.33 From Microsoft SQL Server to Kafka

Supported Source and Destination Databases

Table 4-358 Supported databases

Source DB	Destination DB
RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019)	Kafka 0.11 or later

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-359 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-359 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> • Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER • Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY • Scope of incremental synchronization <ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see **Table 4-360**. DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 4-360 Database account permission

Type	Incremental
Source database user	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized

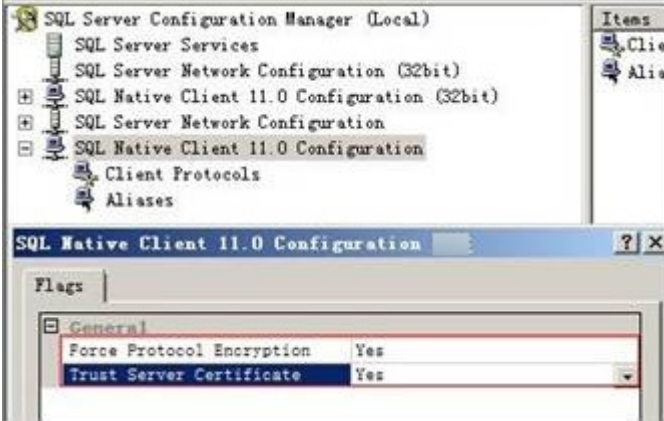
Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-361 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - The SQL Server Agent proxy service must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 4-298. <p style="text-align: center;">Figure 4-298 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database parameter requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later.

Type	Restrictions
	<ul style="list-style-type: none"> - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • DDL operations performed on the source database will not be synchronized to the destination database. • The IMAGE, TEXT, and NTEXT big data types cannot be deleted. • You can add additional synchronization objects.

Procedure

This section uses Microsoft SQL Server to Kafka synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-299 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.

The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

▼

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

▼

* Task Name

DRS-5678
🗑

Description

🗑

0/256

Table 4-362 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 4-300 Synchronization instance details

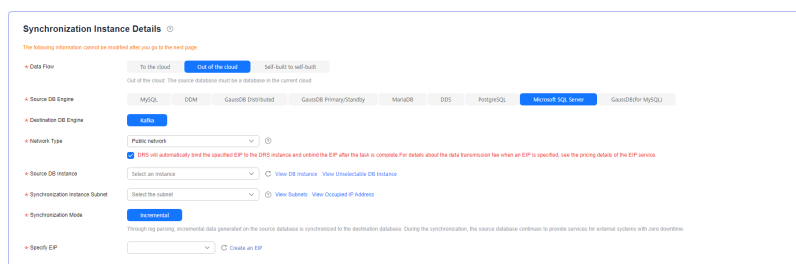


Table 4-363 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select Kafka .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Select an RDS for SQL Server instance you created as the source database.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-301 AZ




Table 4-364 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-302 Enterprise Project and Tags




Table 4-365 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-303 Source database information

Source Database

DB Instance Name

Database Username

Database Password

Table 4-366 Source database settings

Parameter	Description
DB Instance Name	The RDS for SQL Server instance you selected when creating the task. The parameter cannot be changed.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Figure 4-304 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 4-367 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-305 Synchronization Object

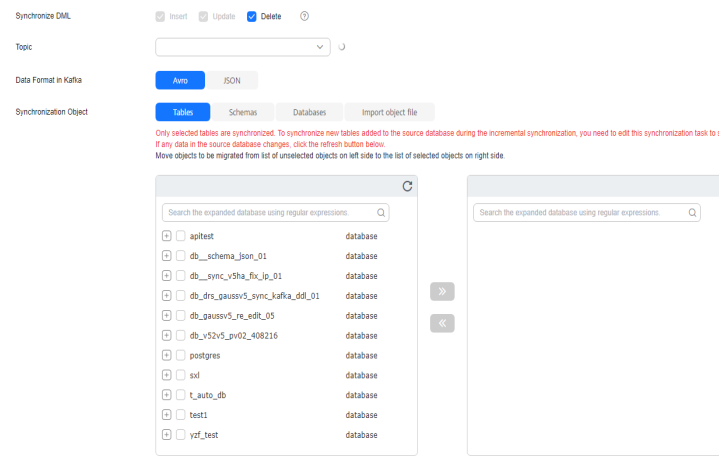



Table 4-368 Synchronization mode and object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Topic	Select the topic to be synchronized to the destination database.

Parameter	Description
Data Format in Kafka	<p>Select the format of data delivered to Kafka.</p> <ul style="list-style-type: none"> ● Avro refers to binary encoded format. ● Json refers to data interchange format. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> ● If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. For details about the import procedure and description, Importing Synchronization Objects. ● When you select Import object file, you can use the mapping function in Changing Object Names (Mapping Object Names) only when the topic synchronization policy is set to A specific topic. Otherwise, topics are generated based on the name format. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-306 Task startup settings





Table 4-369 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.34 From GeminiDB Redis to Redis

Supported Source and Destination Databases

Table 4-370 Supported databases

Source DB	Destination DB
GeminiDB Redis	<ul style="list-style-type: none"> On-premises Codis clusters (open-source Codis 3.0 or later) ECS-hosted Codis clusters (open-source Codis 3.0 or later) On-premises single-node Redis 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x ECS-hosted single-node Redis 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x On-premises master/standby Redis 4.0.x and 5.0.x ECS-hosted master/standby Redis 4.0.x and 5.0.x GeminiDB Redis

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-371 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-371 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• Object level: All• Supported synchronization objects:<ul style="list-style-type: none">- Data of the String, Hash, List, Set, and Sorted Set types can be synchronized.- Lua scripts and transactions cannot be synchronized.- Stream, Exhash, and Bloomfilter data types cannot be synchronized.- User-defined types are not supported.- Commands that are not supported by the destination database cannot be synchronized.

Suggestions

CAUTION

To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.

-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - During the synchronization, ensure that no data is written to the destination database to keep data consistency before and after the synchronization.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full synchronization, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single

full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-372 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> - The source must be a GeminiDB Redis instance on the current cloud. - A GeminiDB Redis instance can be used as the source database of only one DRS task. - The source database GeminiDB Redis kernel version (gemini version) must be 4.2.0 or later. For an incremental task, the source database kernel version (gemini version) must be 6.3.0 or later. - The sequence number of the source database with data stored must be less than the maximum sequence number of the destination database. For example, if the source database contains No.10 database with data stored, the maximum sequence number of the destination database must be greater than or equal to 10. - To create a full+incremental task, the PSYNC command must be supported. - The replication function must be enabled for the source GeminiDB Redis instance. <code>set config set enable-replication 1</code> - To use data migration, disable the configuration item notify-keyspace-events. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database must be empty before the synchronization. - The destination DB instance must have sufficient storage space. - If the destination is master/standby Redis 5.0, enter all IP addresses and ports of the master, slave, and sentinel nodes for the destination database connection testing. The passwords of the master, slave, and sentinel nodes must be the same. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is set on the master and slave nodes, DRS regards the master/standby Redis as a single-node Redis. The master/standby switchover of the destination database is not supported. For the destination database connection testing, enter only the IP address and port of the master node. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is not set on the master or slave node, the master/standby switchover of the destination database is supported. For the destination database connection testing, enter all IP addresses and ports of the master, slave, and sentinel nodes.

Type	Restrictions
	<ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files. - If the replication timeout interval of the source GeminiDB Redis is too short (the default value is 3600s), you can increase the value of repl-timeout to prevent task resuming failure due to a long task failure. - During the synchronization, if the task fails due to configuration increase of the source GeminiDB Redis, you need to reset the task. To ensure data consistency, clear the data that has been migrated to the destination database before reconfiguring the task. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - During the synchronization, the specifications of the source GeminiDB Redis cannot be downgraded. If the specifications of the source database needs to be downgraded, create a synchronization task again. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - For List objects, duplicated data may occur because the existing data on the destination end is not cleared when PSYNC is invoked for retransmission. - If the source is a master/standby GeminiDB Redis instance, commands containing multiple keys (such as mset) cannot be executed in the source database. Otherwise, the error message "CROSSSLOT Keys in request don't hash to the same slot" may be displayed because the keys are not in the same slot or involve multiple slots.
Full synchronization	<ul style="list-style-type: none"> ● During synchronization, do not change the passwords or ports of the source and destination databases. ● During the synchronization, do not modify the destination database.
Incremental synchronization	<ul style="list-style-type: none"> ● During synchronization, do not change the passwords or ports of the source and destination databases. ● During the synchronization, do not modify the destination database.

Prerequisites

- [You have logged in to the DRS console.](#)
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management](#).

Procedure

This section describes how to use DRS to synchronize data from GeminiDB Redis to a Redis database over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-307 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region .. v

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project .. v

* Task Name ⓘ

Description ⓘ

0/256

Table 4-373 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

- Synchronization instance information

Figure 4-308 Synchronization instance information

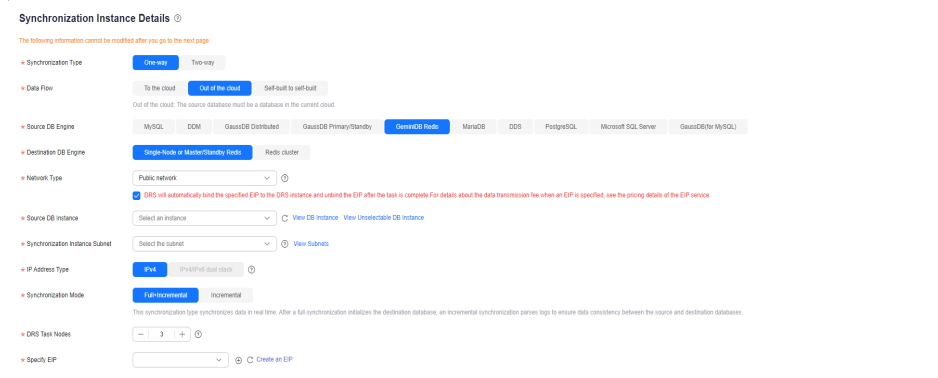


Table 4-374 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select GeminiDB Redis .
Destination DB Engine	Select Redis .
Network Type	Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Select the source GeminiDB Redis instance you have created.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<ul style="list-style-type: none"> - Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database. - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
DRS Task Nodes	<p>Specifies the number of nodes to be migrated using DRS. One DRS task node connects to two to three shards on the source database. If the source GeminiDB Redis database is a master/standby instance, the number of DRS task nodes can only be 1.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-309 AZ




Table 4-375 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-310 Enterprise Project and Tags




Table 4-376 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 4-311 DNS Server

Configure Your Own DNS Server ?

DNS Server

DNS Server IP Address . .

Table 4-377 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use the IP address of your own DNS server as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data migration.

NOTE

This function is available when you need to use the IP address of your own DNS server as the source or destination database IP address.

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 4-312 Source database information

Source Database

DB Instance Name

Database Password

Table 4-378 Source database information

Parameter	Description
DB Instance Name	The GeminiDB Redis instance selected when you create the synchronization task. This parameter cannot be changed.
Database Password	Password of the source database administrator.

 **NOTE**

The IP address or domain name, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Destination database configuration

Figure 4-313 Destination database information

Destination Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Password ⓘ

Table 4-379 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <ul style="list-style-type: none"> - If a single-node or proxy instance is used, ensure that the entered IP addresses or domain names belong to the same instance. - If the destination is master/standby Redis 5.0, enter all IP addresses and ports of the master, slave, and sentinel nodes for the destination database connection testing. The passwords of the master, slave, and sentinel nodes must be the same. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is set on the master and slave nodes, DRS regards the master/standby Redis as a single-node Redis. The master/standby switchover of the destination database is not supported. For the destination database connection testing, enter only the IP address and port of the master node. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is not set on the master or slave node, the master/standby switchover of the destination database is supported. For the destination database connection testing, enter all IP addresses and ports of the master, slave, and sentinel nodes.
Database Password	The password for the database username.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-314 Synchronization objects




Table 4-380 Synchronization Object

Parameter	Description
Synchronization Object	All database objects can be synchronized.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, and click **Submit** to submit the task.

Figure 4-315 Task startup settings




Table 4-381 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. The Start at a specified time option is recommended. NOTE The synchronization task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the task billing is about to start, the status, latency metric, or data of the synchronization task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none">In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.Before setting the delay threshold, enable Send Notifications.If the delay threshold is set to 0, no notifications will be sent to the recipient.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4.35 From GeminiDB Redis to Redis Cluster

Supported Source and Destination Databases

Table 4-382 Supported databases

Source DB	Destination DB
GeminiDB Redis	<ul style="list-style-type: none"> • ECS-hosted Redis Cluster (open-source Redis 4.0 or later) • On-premises Redis Cluster (open-source Redis 4.0 or later)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-383 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-383 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> • Object level: All • Supported synchronization objects: <ul style="list-style-type: none"> - Data of the String, Hash, List, Set, and Sorted Set types can be synchronized. - Lua scripts and transactions cannot be synchronized. - Stream, Exhash, and Bloomfilter data types cannot be synchronized. - User-defined types are not supported. - Commands that are not supported by the destination database cannot be synchronized. - Batch operation commands without hash_tag cannot be synchronized.

Suggestions

⚠ CAUTION

To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- During the synchronization, ensure that no data is written to the destination database to keep data consistency before and after the synchronization.
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full synchronization, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-384 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> - The source must be a GeminiDB Redis instance on the current cloud. - A GeminiDB Redis instance can be used as the source database of only one DRS task. - The source GeminiDB Redis instance kernel version (gemini version) must be 4.2.0 or later. - To use data synchronization, disable the configuration item notify-keyspace-events. - To create a full+incremental task, the PSYNC command must be supported. - The replication function must be enabled for the source GeminiDB Redis instance. <pre>set config set enable-replication 1</pre> ● Destination database requirements: <ul style="list-style-type: none"> - The destination database must be empty before the synchronization. - The destination DB instance must have sufficient storage space. - The username cannot be entered in Redis 6.0. - If the destination cluster instance is created using containers, ECSs, or NAT, ensure that the config settings of the cluster allow DRS to discover all master and standby nodes. ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files. - If the replication timeout interval of the source GeminiDB Redis is too short (the default value is 3600s), you can increase the value of repl-timeout to prevent task resuming failure due to a long task failure. - During the synchronization, if the task fails due to configuration increase of the source GeminiDB Redis, you need to reset the task. To ensure data consistency, clear the data that has been migrated to the destination database before reconfiguring the task. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - During the synchronization, the specifications of the source GeminiDB Redis cannot be downgraded. If the specifications of the source database needs to be downgraded, create a synchronization task again. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.)

Type	Constraints
	<ul style="list-style-type: none"> - For List objects, duplicated data may occur because the existing data on the destination end is not cleared when PSYNC is invoked for retransmission. - If the source is a master/standby GeminiDB Redis instance, commands containing multiple keys (such as mset) cannot be executed in the source database. Otherwise, the error message "CROSSSLOT Keys in request don't hash to the same slot" may be displayed because the keys are not in the same slot or involve multiple slots.
Full synchronization	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During the synchronization, do not modify the destination database.
Incremental synchronization	<ul style="list-style-type: none"> • During synchronization, do not change the passwords or ports of the source and destination databases. • During the synchronization, do not modify the destination database.

Prerequisites

- [You have logged in to the DRS console.](#)
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management](#).

Procedure

This section describes how to use DRS to synchronize data from GeminiDB Redis to Redis Cluster over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-316 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 4-385 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\\"

- Synchronization instance information

Figure 4-317 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

* Synchronization Type Two-way

* Data Flow
Out of the cloud: The source database must be a database in the current cloud.

* Source DB Engine MySQL DDM GaussDB Distributed GaussDB Primary/Standby MariaDB DDS PostgreSQL Microsoft SQL Server GaussDB for MySQL

* Destination DB Engine

* Network Type ⓘ
 DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

* Source DB Instance ⓘ [View DB Instance](#) [View Unavailable DB Instance](#)

* Synchronization Instance Subnet ⓘ [View Subnets](#)

* IP Address Type ⓘ

* Synchronization Mode Incremental
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

* DRS Task Nodes ⓘ

* Specify EIP ⓘ [Create an EIP](#)

Table 4-386 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select GeminiDB Redis .
Destination DB Engine	Select Redis cluster .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Select the source GeminiDB Redis instance you have created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database. - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p>
DRS Task Nodes	Specifies the number of nodes to be migrated using DRS. One DRS task node connects to two to three shards on the source database. If the source GeminiDB Redis database is a master/standby instance, the number of DRS task nodes can only be 1.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-318 AZ

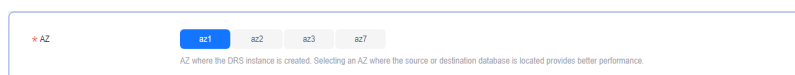


Table 4-387 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 4-319 Enterprise Project and Tags




Table 4-388 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Configuring your own DNS server

Figure 4-320 DNS Server

Configure Your Own DNS Server 

DNS Server

DNS Server IP Address

Table 4-389 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use your own private domain name as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data synchronization.

 **NOTE**

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 4-321 Source database information

Source Database

DB Instance Name

Database Password

Table 4-390 Source database settings

Parameter	Description
DB Instance Name	The GeminiDB Redis instance selected when you created the migration task. This parameter cannot be changed.
Database Password	Password of the source database administrator.

 **NOTE**

The IP address or domain name, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

- Destination database configuration

Figure 4-322 Destination database information

Destination Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Password ⓘ

Table 4-391 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. Enter the IP addresses and ports of all master and slave nodes in the Redis Cluster instance. A maximum of 32 IP addresses or domain names can be entered. Use commas (,) to separate multiple values. For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Password	The password for the database username.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the password are permanently deleted.

- Step 4** On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-323 Synchronization objects

Synchronization Object



Table 4-392 Synchronization Object

Parameter	Description
Synchronization Object	All database objects can be synchronized.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, and click **Submit** to submit the task.

Figure 4-324 Task startup settings

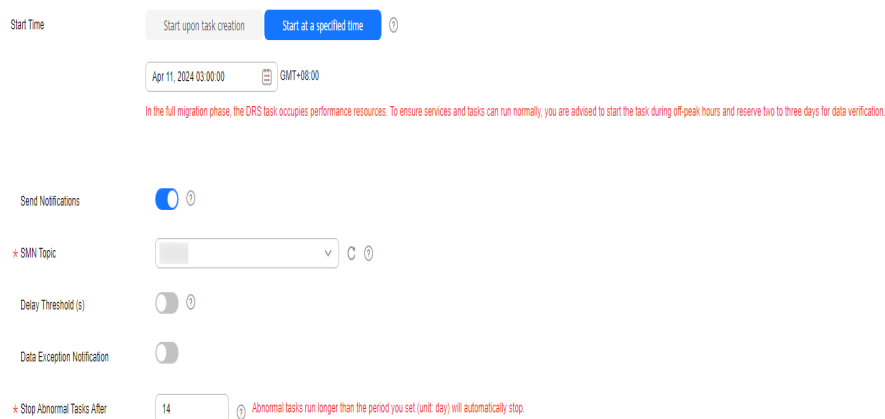



Table 4-393 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. The Start at a specified time option is recommended. NOTE The synchronization task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the task billing is about to start, the status, latency metric, or data of the synchronization task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none">• In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.• Before setting the delay threshold, enable Send Notifications.• If the delay threshold is set to 0, no notifications will be sent to the recipient.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5 Between Self-built Databases

5.1 From MySQL to Kafka

Supported Source and Destination Databases

Table 5-1 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECS	<ul style="list-style-type: none">Kafka

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- Tables to be synchronized without a primary key may be locked for 3s.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-2 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, SHOW VIEW, EVENT, REPLICATION SLAVE, and REPLICATION CLIENT. For a full +incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.
Synchronization object	<ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. <ul style="list-style-type: none"> - If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The database and table names in the source database cannot contain non-ASCII characters, or the following special characters: <code>.'<'>\/\"</code> ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.
Destination database	<ul style="list-style-type: none"> ● The destination database is a Kafka database. ● You are advised to set auto.create.topics.enable of Kafka to false.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● During the synchronization, do not delete or change the username, password, or permission of the source database, or change the port of the destination database. ● Data inconsistency may occur when the MyISAM table is modified during synchronization. ● During synchronization of table-level objects, renaming tables is not recommended. ● During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-1 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page. The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 5-3 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-2 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

To the cloud
Out of the cloud
Default is self-host

MySQL
Oracle
DB2 for LUW
GreenDB Distributed
GreenDB Primary/Standby
PostgreSQL

CDB&SO
GreenDB Primary/Standby
Auto

Public network
ⓘ

DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

Single-AZ
Dual-AZ

Single-AZ deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service availability.

VPC
New VPC

Synchronization Instance Subnet
New Subnets
View Occupied IP Address

Security Group
Default
ⓘ

Full-consistent
Incremental

The synchronous full synchronization data in real time. After a full synchronization releases the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Security EP
Create an EP

Table 5-4 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select MySQL .
Destination DB Engine	Select Kafka .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
DRS Task Type	Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task .
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-3 Task type





Table 5-5 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-4 AZ</p> 

- Enterprise Project and Tags

Figure 5-5 Enterprise Project and Tags




Table 5-6 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-6 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-7 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-7 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol v ⓘ

● Test successful

Table 5-8 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

The parameters on the **Set Synchronization Task** page vary according to the synchronization mode selected in **Step 2**. The **Full+Incremental** synchronization mode is used as an example in **Figure 5-8**.

Figure 5-8 Synchronization mode

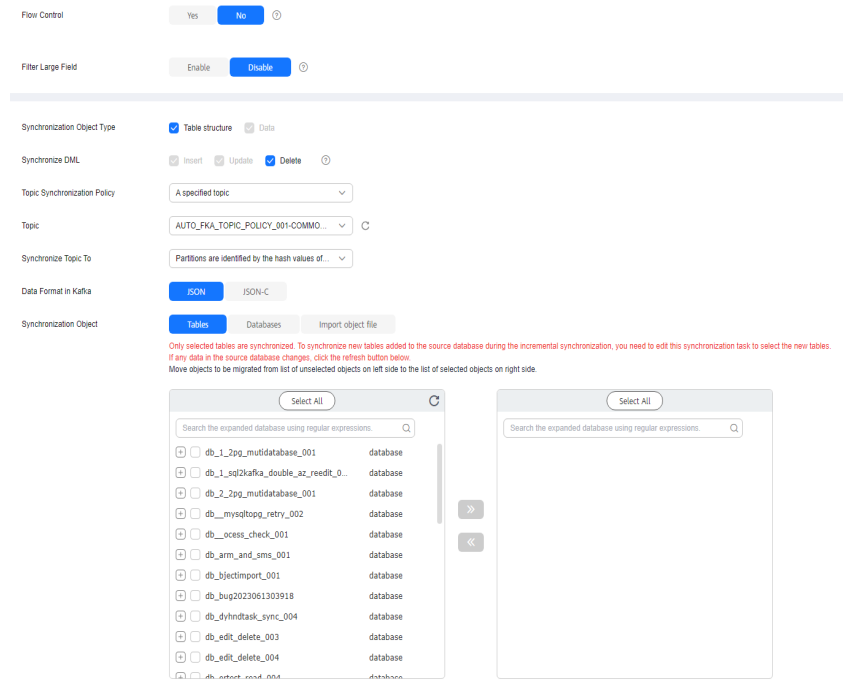
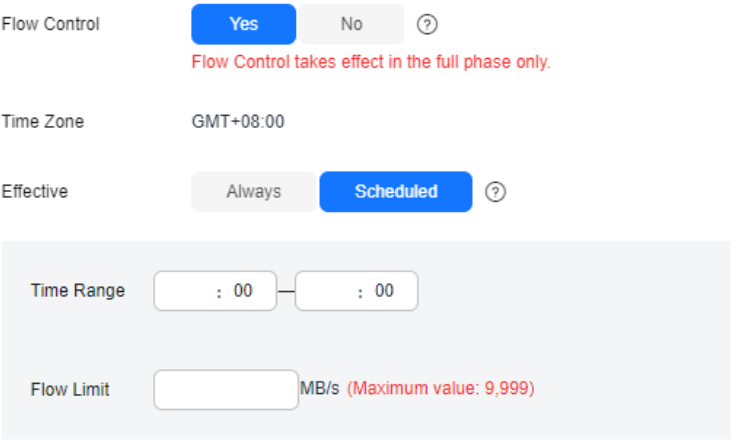
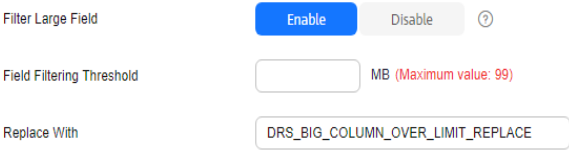



Table 5-9 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-9 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
<p>Filter Large Field</p>	<p>Indicates whether to use large field filtering to process special fields (blob, mediumblob, longblob, varbinary, mediumtext and longtext) in a synchronization table.</p> <ul style="list-style-type: none"> <p>Enable</p> <p>You need to set Field Filtering Threshold and Replace With. If the size of a field exceeds the threshold, the value is replaced based on a specified character.</p> <p>Note that large field filtering is used to replace the value of a field, not the entire DML record. If a DML record contains many large fields, the size of only some of these fields exceeds the filtering threshold, and the accumulated value of other fields that do not exceed the filtering threshold is greater than the value of request.max.size, when data is written to Kafka, the size of the message body in the destination Kafka database may still exceed the upper limit, resulting in a DRS error.</p> <p>Figure 5-10 Setting large field filtering</p>  <p>Disable</p> <p>Large fields are not filtered.</p>
<p>Synchronization Object Type</p>	<p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p>
<p>Synchronize DML</p>	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
<p>Start Point</p>	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name.</p> <p>For example, if this parameter is set to \$database\$-tablename\$ and the database name is db1, and the table name is tab1, the topic name is db1-tab1. If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>

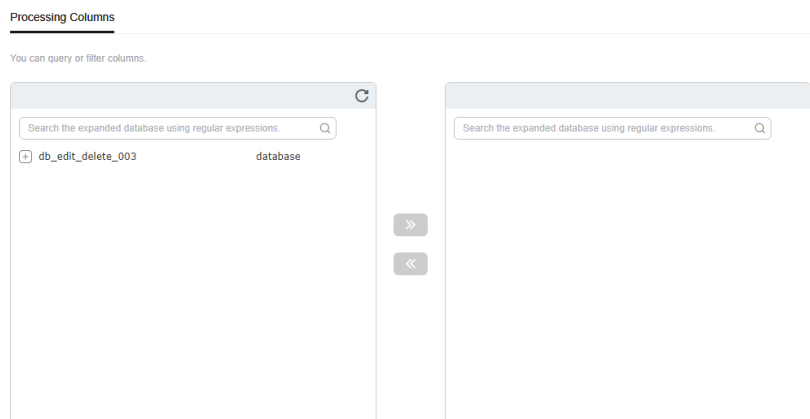
Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • Partitions are identified by the hash values of the primary key: This mode applies to scenarios where a single table contains a large amount of data, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. Data sequence can be preserved only when the primary key value is not changed. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.table_name.
Data Format in Kafka	<p>Select the data format to be delivered from MySQL to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <ul style="list-style-type: none"> – When importing an object file, you can perform either topic mapping or object name mapping. – If you perform topic mapping when importing an object file, different tables can be synchronized to different topics in the destination database. If topic mapping is not specified for an object, the object uses the external topic policy. You can modify the mapping when editing the synchronization object. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 5-11 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-12 Task startup settings

Table 5-10 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.2 From MySQL to CSS/ES

Supported Source and Destination Databases

Table 5-11 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MySQL database (versions 5.5, 5.6, 5.7, and 8.0)Self-built MySQL database on ECSs (versions 5.5, 5.6, 5.7, and 8.0)	<ul style="list-style-type: none">ElasticSearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9 and 7.10

 NOTE

Only whitelisted users can use this function.

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-12 Precautions

Type	Restrictions
Database permissions	Minimum permission requirements for full plus incremental synchronization: <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the following permissions: READ and WRITE

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> • The table data can be synchronized. • Databases, views, indexes, constraints, functions, stored procedures, triggers, and events cannot be synchronized. • The system database and event status cannot be synchronized. • Tables whose primary keys are of the FLOAT type cannot be synchronized. • Tables that do not have primary keys cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>\'\" ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>\'\" ● The column name in the source database cannot contain non-ASCII characters or the following characters: '\." ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. <ul style="list-style-type: none"> - If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.

Type	Restrictions
	<ul style="list-style-type: none">• The source database does not support the mysql binlog dump command.• The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails.• The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.• The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail.
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly.• The destination DB instance must have sufficient storage space.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored to a point in time when a full synchronization was being performed. ● Convert the value range of the source database time field that is not supported by the destination database to null. ● The strings padded with zeros in the source database may be truncated because the source database uses the fixed-length binary data type, and the destination database uses the variable-length data type. ● If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The binary value is encrypted using Base64 and then written to the destination database. ● If no time zone is specified for the source database, specify the time zone for synchronizing the datetime type to the destination database. ● All table field names are converted to lowercase letters. ● If the <code>_id</code> field of the destination database is generated using multiple columns in the source database, separate these columns with colons (:). ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.

Type	Restrictions
	<ul style="list-style-type: none"> • During the synchronization, do not create a database named ib_logfile in the source. • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable data transfer is supported. If the host system breaks down, data may be repeatedly inserted into non-transactional tables that do not have primary keys. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-13 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page.
The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256

Table 5-13 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-14 Synchronization instance details

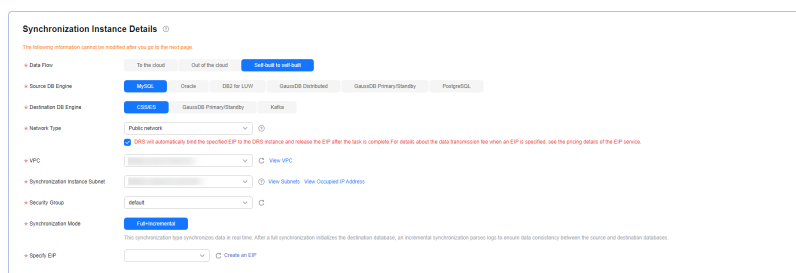


Table 5-14 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select MySQL .
Destination DB Engine	Select CSS/ES .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
VPC	Select an available VPC.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental , data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-15 Task type




Table 5-15 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-16 Enterprise Project and Tags




Table 5-16 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-17 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-17 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

Parameter	Description
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-18 Destination database information

Destination Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

SSL Connection

This button is available only after the replication instance is created successfully.

Table 5-18 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. Only .cer and .pem certificates are supported. If SSL is disabled, your data may be at risk.

 NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 5-19 Synchronization mode

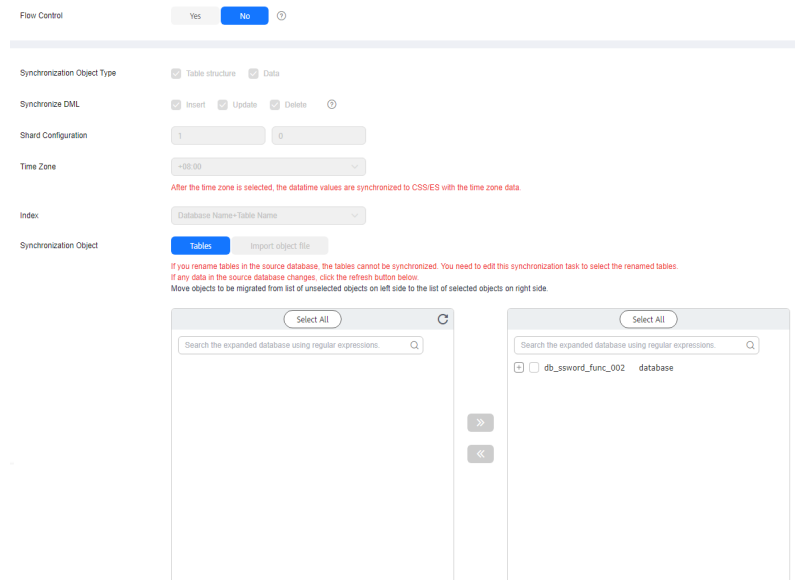
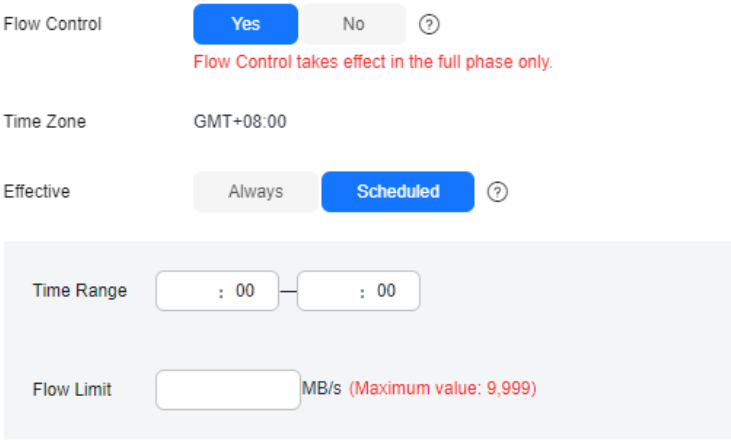



Table 5-19 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-20 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

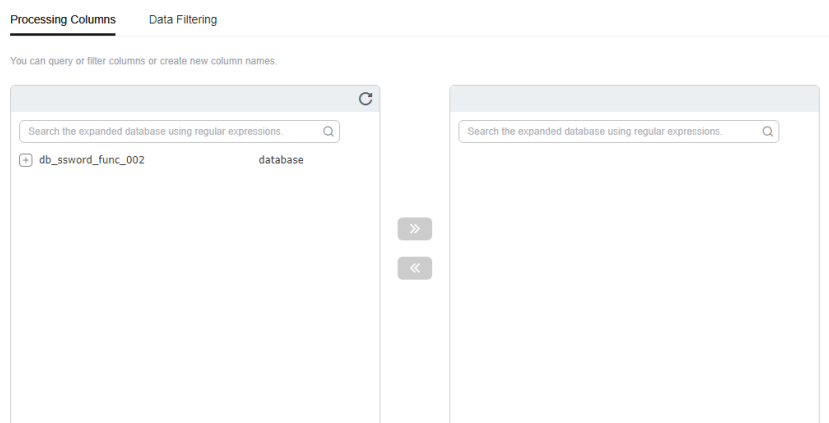
Parameter	Description
Synchronization Object Type	<p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Shard Configuration	<p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p>
Time Zone	<p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p>
Index Name	<ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> You can select objects by importing a file. For details, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). If you select Tables, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the limit is exceeded, you can add synchronization objects in batches by editing the synchronization objects. For details, see Editing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-21 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-22 Task startup settings

Table 5-20 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.3 From MySQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-21 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS 	GaussDB Primary/Standby (B-compatible or M-compatible)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-22](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-22 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	SELECT Statement: GRANT SELECT ON <database>.<table> to <user>;	SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>; For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user		<ul style="list-style-type: none"> • Database-level permissions: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> • Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> • Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. • If the destination database is in M-compatible mode, the Sysadmin role must be configured. Authorization statement: ALTER USER <user> SYSADMIN;

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-23](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-23 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• The names of the source databases, tables, and columns cannot contain non-ASCII characters or the following characters: <'>/\• The column names in the source database tables cannot end with a backslash (\).• Only table structures, table data, and indexes can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed.• Other database objects such as stored procedures cannot be synchronized.• MySQL tables containing virtual columns cannot be synchronized.• Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 5-24](#).

Table 5-24 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: <'>/\"• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • The source database cannot be restored. • During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> • When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. • During table structure synchronization of full synchronization, the auto-increment attribute of the table structure is not synchronized for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames and permissions of the source and destination databases or change the ports of the source and destination databases. • If the session variable character_set_client is set to binary, some data may include garbled characters. • During incremental synchronization, the following DDL operations are supported by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, RENAME_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, ALTER_COLUMN, DROP_INDEX, and RENAME_INDEX. You can select the DDL operations to be synchronized on the object selection page as required. • During an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated on the destination database. <p>Stopping a task</p>

Type	Restrictions
	<ul style="list-style-type: none">● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none">● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● By default, the MySQL-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● After a table in the source database is synchronized to the destination database, the table is distributed in hash mode and cannot be replicated. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. – Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The changes to the source database cannot be synchronized to the destination database in multiple tasks at the same time. Otherwise, data inconsistency and synchronization failure may occur.

Type	Restrictions
	<ul style="list-style-type: none"> ● Only the database-level character set can be set for GaussDB. If the table-level and field-level character sets are set for MySQL, the table structure may fail to be created due to character length differences. ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of columns in the source database. For example, if the character set of the source database is UTF8, increase the length (byte) by three times. If the character set of the source database is UTF8MB4, increase the length (byte) by four times. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. ● If a time type is used as a primary key and the value of the time type is an invalid value of the destination database, data inconsistency or task failure may occur during incremental synchronization. ● For incremental DDL synchronization, the DDL syntax in B compatibility mode is used for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). As a result, incremental synchronization may fail. ● Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. ● If a table in the source MySQL database contains a binary field with a fixed length, the MySQL driver adds \0 to the end of the data based on the length. As a result, there may be data inconsistency after the data is synchronized to the destination GaussDB database. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot

Type	Restrictions
	<p>be written in strict mode, there may be data inconsistency during synchronization.</p> <ul style="list-style-type: none"> ● If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000. ● If the table structure is synchronized, column names in the destination database will be converted to lowercase letters by default. If the table to be synchronized already exists in the destination database, the column names of the synchronization object table in the destination database must be in lowercase. If column names in the destination database are required to be case sensitive, choose Service Tickets > Create Service Ticket in the upper right corner of the management console to submit a service ticket. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the time type is used as a primary key and there is abnormal data (data beyond the range from 00:00:00 to 23:59:59), the value comparison result may be inconsistent.

Procedure

This section uses real-time synchronization from MySQL to GaussDB primary/standby as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-23 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name:

Description:

0/256

Table 5-25 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 5-24 Synchronization instance information

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud Out of the cloud Self-built to self-built

Source DB Engine: MySQL Oracle DB2 for LUW GaussDB Distributed GaussDB Primary/Standby PostgreSQL

Destination DB Engine: CDB3 GaussDB Primary/Standby Kafka

Network Type: Public network VPC

VPC:

Synchronization Instance Subnet:

Security Group:

Synchronization Mode: Full Incremental Full Incremental

Specify EP:

Table 5-26 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB Primary/Standby .

Parameter	Description
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-25 Task type




Table 5-27 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-26 Enterprise Project and Tags




Table 5-28 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-27 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 5-29 Source database settings


Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-28 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-30 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 5-29 Synchronization Mode

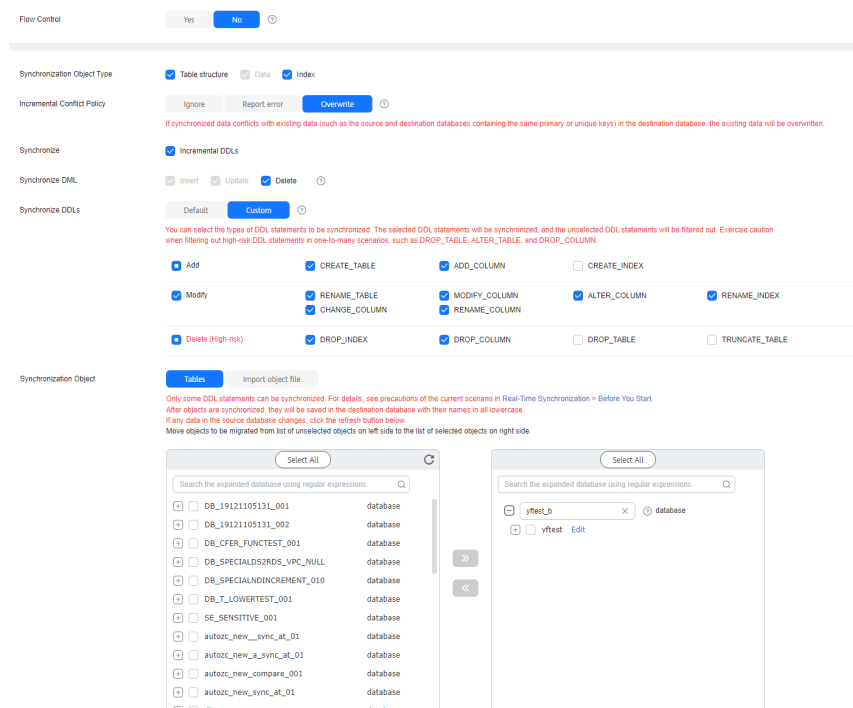
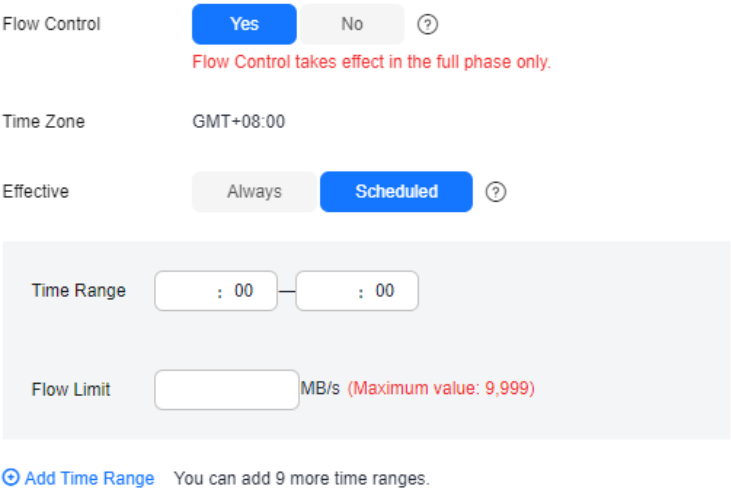



Table 5-31 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-30 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten.
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> – During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom <p>You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.</p> <p>If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-31 Processing data

Belonged Database Table	New Name	Column Name	Operation Type	Type	Operation
gbest_order_payment_1	gbest_order_payment_1	c1	The serverName@database@table column is used	varchar(19)	Add Detail
gbest_order_payment_10	gbest_order_payment_10	c2	Default 1234	int	Add Detail
gbest_order_payment_11	gbest_order_payment_11	--	--	--	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-32 Task startup settings

Start Time Start upon task creation Start at a specified time ⓘ

Send Notifications ⓘ

SMN Topic [dropdown] ⓘ

Delay Threshold (s) ⓘ

Data Exception Notification ⓘ


Stop Abnormal Tasks After [14] ⓘ Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 5-32 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.4 From MySQL to GaussDB Distributed

Supported Source and Destination Databases

Table 5-33 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises MySQL databases • MySQL databases on an ECS 	GaussDB Distributed (MySQL-compatible)

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-34](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-34 Database account permission

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<p>SELECT</p> <p>Statement: GRANT SELECT ON <database>.<table> to <user>;</p>	<p>SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>;</p> <p>For a full+incremental task, if the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user	<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-35](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-35 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none">• The names of the source databases, tables, and columns cannot contain non-ASCII characters or the following characters: <'>/\• The column names in the source database tables cannot end with a backslash (\).• Only table structures, table data, and indexes can be synchronized.• Only tables with primary keys can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed.• Other database objects such as stored procedures cannot be synchronized.• A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.• Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 5-36](#).

Table 5-36 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"><li data-bbox="587 344 1426 443">• The names of the source databases and tables cannot contain non-ASCII characters, or the following special characters: .<'>/"<li data-bbox="587 454 1426 553">• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.<li data-bbox="587 564 1426 698">• If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> Do not perform data restoration on the source database. During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. During table structure synchronization of full synchronization, the auto-increment attribute of the table structure is not synchronized for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. <p>Incremental synchronization</p> <ul style="list-style-type: none"> Do not modify or delete the usernames and permissions of the source and destination databases or change the ports of the source and destination databases. If the session variable character_set_client is set to binary, some data may include garbled characters. During incremental synchronization, the following DDL operations are supported by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, RENAME_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, ALTER_COLUMN, DROP_INDEX, and RENAME_INDEX. You can select the DDL operations to be synchronized on the object selection page as required. During an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated on the destination database. <p>Stopping a task</p>

Type	Restrictions
	<ul style="list-style-type: none">● Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.● Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none">● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● By default, the MySQL-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● DRS does not proactively specify a distribution column and distribution mode, which is determined by the database itself. ● Replication tables without primary keys cannot be synchronized. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. ● The changes to the source database cannot be synchronized to the destination database in multiple tasks at the same

Type	Restrictions
	<p>time. Otherwise, data inconsistency and synchronization failure may occur.</p> <ul style="list-style-type: none"> ● Only the database-level character set can be set for GaussDB. If the table-level and field-level character sets are set for MySQL, the table structure may fail to be created due to character length differences. ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of columns in the source database. For example, if the character set of the source database is UTF8, increase the length (byte) by three times. If the character set of the source database is UTF8MB4, increase the length (byte) by four times. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. ● If a time type is used as a primary key and the value of the time type is an invalid value of the destination database, data inconsistency or task failure may occur during incremental synchronization. ● For incremental DDL synchronization, the DDL syntax in B compatibility mode is used for data synchronization from MySQL to GaussDB (PostgreSQL-compatible). As a result, incremental synchronization may fail. ● Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. ● If a table in the source MySQL database contains a binary field with a fixed length, the MySQL driver adds \0 to the end of the data based on the length. As a result, there may be data inconsistency after the data is synchronized to the destination GaussDB database.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000. ● If the table structure is synchronized, column names in the destination database will be converted to lowercase letters by default. If the table to be synchronized already exists in the destination database, the column names of the synchronization object table in the destination database must be in lowercase. If column names in the destination database are required to be case sensitive, choose Service Tickets > Create Service Ticket in the upper right corner of the management console to submit a service ticket. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If the time type is used as a primary key and there is abnormal data (data beyond the range from 00:00:00 to 23:59:59), the value comparison result may be inconsistent.

Procedure

This section uses real-time synchronization from self-managed MySQL to GaussDB distributed as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-33 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-37 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-34 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud Self-built to self-built

• Source DB Engine MySQL Oracle DB for LUW GaussDB Distributed GaussDB Primary/Standby PostgreSQL

• Destination DB Engine COSDB GaussDB Standalone GaussDB Primary/Standby Kafka

• Network Type Public network VPC
DRS will automatically bind the specified EP to the DRS instance and unbind the EP after the task is complete. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

• VPC

• Synchronization Instance Subnet

• Security Group

• Synchronization Mode Full-incremental Full Incremental
This synchronization task synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization passes logs to ensure data consistency between the source and destination databases.

• Specify EP

Table 5-38 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select MySQL .

Parameter	Description
Destination DB Engine	Select GaussDB Distributed .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-35 Task type




Table 5-39 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-36 Enterprise Project and Tags




Table 5-40 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-37 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-41 Source database settings


Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-38 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-42 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 5-39 Synchronization Mode

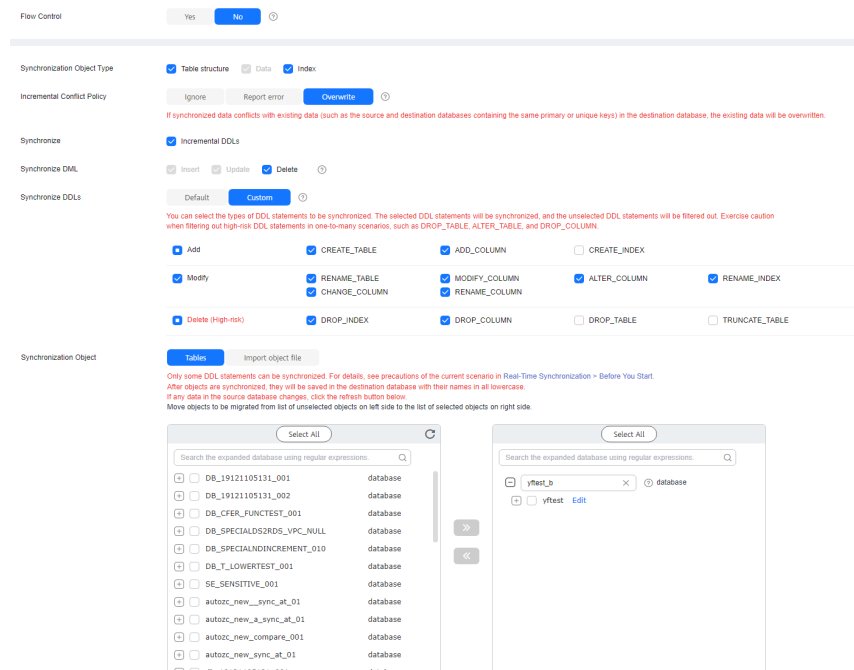
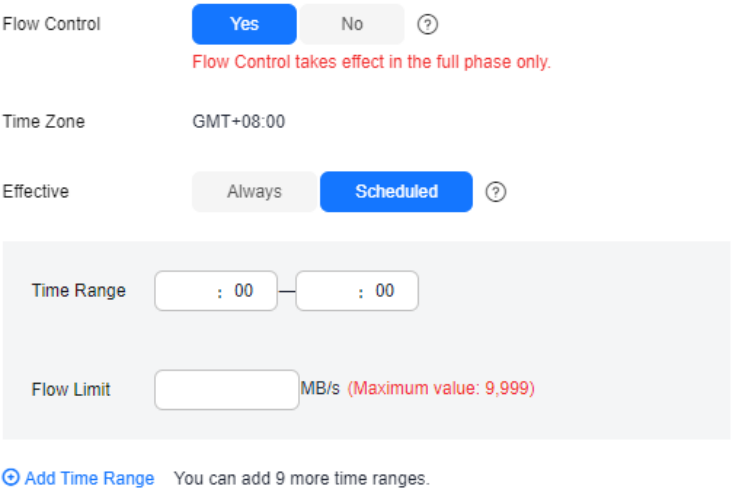



Table 5-43 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-40 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten.
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> – During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are <code>ADD_COLUMN</code>, <code>MODIFY_COLUMN</code>, and <code>ALTER_COLUMN</code>. ● Custom <p>You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.</p> <p>If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-41 Processing data

Belonged Database Table	New name	Column Name	Operation Type	Type	Operation
test-database_000	test_RS_newdatabase_000	-	-	-	Add
test-mysql1	test_RS_newtest1_table	-	-	-	Add
test-mysql2	test_RS_newmysql2	-	-	-	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-42 Task startup settings


* Start Time: Start upon task creation Start at a specified time ?
 Send Notifications: ?
 * SMN Topic: ?
 Delay Threshold (s): ?
 Data Exception Notification: ?
 * Stop Abnormal Tasks After: ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 5-44 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.

Parameter	Description
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.5 From Oracle to Kafka

Supported Source and Destination Databases

Table 5-45 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)• Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	<ul style="list-style-type: none">• Kafka

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

Before creating a synchronization task, read the following notes:

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-46 Environment Constraints

Type	Constraint
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. - Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.

Type	Constraint
Synchronization object	<ul style="list-style-type: none"> ● Only table data can be synchronized in real time. ● The following data types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, CHAR, NCHAR, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE, and TIMESTAMP WITH LOCAL TIME ZONE. ● The following column types cannot pass the precheck: GEOMETRY and self-defined. ● The following column types cannot be synchronized but can pass the precheck: INTERVAL_YEAR_TO_MONTH, INTERVAL_DAY_TO_SECOND, UROWID, BFILE and XML. ● The following column types are deleted by default before synchronization: RAW, BLOB, CLOB, NCLOB, XML, LONG and LONG RAW. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute and does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB and XML types cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB or XML type, the incremental synchronization will fail. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● Temporary tables in the source database cannot be synchronized. ● Tables whose default values contain expression functions cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database.

Type	Constraint
Source database	<ul style="list-style-type: none"> • The names of databases and tables cannot contain non-ASCII characters or special characters .><`\ ,?!" • An empty source database cannot be synchronized. • If the source database is an RAC database, you cannot add or delete nodes. • If the source database is an RAC database and uses SCAN IP, the synchronization instance must be able to connect to the virtual IP addresses of all RAC nodes. Otherwise, the connection check fails. • Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.
Destination database	<ul style="list-style-type: none"> • The destination database is a Kafka database. • You are advised to set auto.create.topics.enable of Kafka to false.

Type	Constraint
Precautions	<ul style="list-style-type: none"> ● If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. ● After data in the Oracle database is synchronized to Kafka, the character set becomes UTF8. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● The supplemental log supports all or primary key+unique index columns. ● If a column that is not displayed in the log, it will not be displayed in the transferred message, which means that the column is not updated. ● During synchronization, do not delete the username, password, and permissions of the source and destination databases or modify the port of the destination database.

Type	Constraint
	<ul style="list-style-type: none"> ● During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. ● During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. ● During the synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● In an incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During an incremental synchronization of table-level objects, renaming tables is not recommended. ● If you select Tables for Synchronization Object, all tables must be synchronized to the same topic at the destination end. If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. ● DDL operations can be performed on tables. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. ● Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-43 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page. The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region .. v

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-47 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-44 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

+ Data Flow To the cloud Out of the cloud Sub built to self built

+ Source DB Engine MySQL Oracle DB2 for LUW GaussDB Distributed GaussDB Primary/Standby PostgreSQL

+ Destination DB Engine GaussDB Distributed GaussDB Primary/Standby Kafka

+ Network Type Public network ⓘ

+ DRS Task Type Single-AZ Dual-AZ

Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime.

+ VPC vpc-1 ⓘ View VPC

+ Synchronization Instance Subnet subnet-1 ⓘ View Subnets View Occupied IP Address

+ Security Group default ⓘ

+ Synchronization Mode Incremental

Through IP learning, incremental data generated on the source database is synchronized to the destination database. During the synchronization, the source database continues to provide services for external systems with zero downtime.

+ Specify EIP eip-1 ⓘ Create an EIP

Table 5-48 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select Oracle .
Destination DB Engine	Select Kafka .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-45 Task type





Table 5-49 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-46 AZ</p> 

- Enterprise Project and Tags

Figure 5-47 Enterprise Project and Tags




Table 5-50 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-48 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

PDB Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

This button is available only after the replication instance is created successfully.

Table 5-51 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 5-49 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 5-52 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.

Parameter	Description
Security Protocol	Available options: PLAINTEXT, SSL, SASL_PLAINTEXT, and SASL_SSL. For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select a topic and objects to be synchronized, and then click **Next**.

Figure 5-50 Synchronization mode

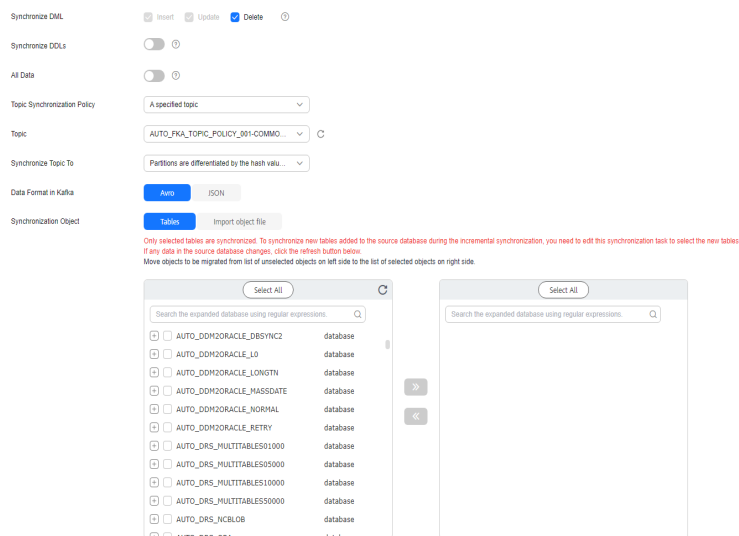



Table 5-53 Synchronization mode and object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Synchronize DDLs	Controls whether to synchronize DDLs to Kafka. If Synchronize DDLs is enabled and Partitions are identified by the hash values of the primary key is selected, DDLs are hashed based on the table name because the DDLs do not have the primary key value. In other cases, the synchronization policy is the same as the partition policy.

Parameter	Description
All Data	<p>Controls whether to synchronize all data in a single row. DRS parses the source database logs to synchronize incremental data. The data integrity in a single row depends on whether the values of all columns are recorded in the logs.</p> <p>If all data is required for the synchronization object, all-level supplemental logging must be enabled in the source database to record all column values of a single row data. This option is associated with the verification of the supplemental logging level in the source database in the pre-check phase. DRS incremental synchronization has the minimum requirement of table-level PK/UI supplemental logging. For details, see How Do I Check Supplemental Logging of the Source Oracle Database?</p>
Topic Synchronization Policy	<p>Topic synchronization policy. The options are as follows:</p> <ul style="list-style-type: none"> • Select A specified topic if the data volume of the source database is small. • Select Automatically generated based on the schema name if each schema contains a lot of data. • Select Automatically generated using the schema_name-table_name format if each table contains a lot of data.
Topic	<p>Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy is set to A specified topic.</p>
Topic Name Format	<p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>The topic name format supports the schema and tablename variables. Other characters are used as constants. Replace \$schema\$ with the schema name and \$tablename\$ with the table name. For example, if this parameter is set to \$schema\$-\$tablename\$, the schema name is schema1, and the table name is tab1 when Oracle is the source, the topic name is schema1-tab1.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>

Parameter	Description
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by the hash values of <code>schema_name.table_name</code>, the performance on a single table query can be improved. • If topics are synchronized to different partitions by the hash values of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of <code>schema_name.table_name</code>. • Partitions are differentiated by the hash values of schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing data in multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, data is sent using multiple threads by default. This ensures strong consistency but write performance is impacted. If strong transaction consistency is required, you are advised to select this option and contact O&M personnel to change to single-thread Kafka write, or set the topic synchronization policy to Automatically generated based on the table name.
Data Format in Kafka	<p>Select the format of data sent from the Oracle database to the Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. • Json refers to data interchange format. <p>For details, see Kafka Message Format.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. For details about the import procedure and description, Importing Synchronization Objects. • When you select Import object file, you can use the mapping function in Changing Object Names (Mapping Object Names) only when the topic synchronization policy is set to A specific topic. Otherwise, topics are generated based on the name format. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-51 Task startup settings

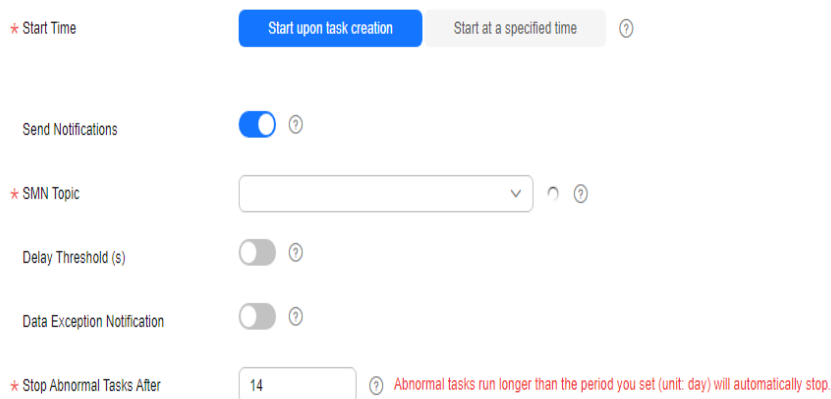



Table 5-54 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.6 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-55 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) 	GaussDB primary/standby

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization

tasks require different permissions. For details, see [Table 5-56](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-56 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user		<ul style="list-style-type: none"> • Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> • Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> • Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-57](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-57 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● Full synchronization does not support the following column types: bfile, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, sdo_geometry, urowid, interval, and user-defined types. ● The maximum precision supported by timestamp and interval day to second is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute and does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 5-58](#).

Table 5-58 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"><li data-bbox="587 342 1385 472">• The total index length of columns in the source database cannot exceed the length limit in the destination database. For detailed length requirements, see Index Length Description.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> • Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. • Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. • During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize

Type	Restrictions
	<p>the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During table structure synchronization in a full synchronization, only default value constraints of the character string and number types are supported. Default value constraints of the function and sequence types are not supported. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. - During incremental synchronization, tables cannot be renamed, tables cannot be redefined online, and partitions cannot be swapped. ● During incremental synchronization, modify the start point of a capture task to resynchronize data. <ul style="list-style-type: none"> - After the capture point is modified, the previous object-level comparison result is not displayed. - If the start point of a capture task is changed separately, the change is synchronized to the start point of the replay task. That is, the start point of the replay task is the same as that of the capture task. This does not affect the change of the start point of the replay task. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not

Type	Restrictions
	<p>generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency.</p> <ul style="list-style-type: none"> • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. • If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. • For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. • In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: decode, nvl, nvl2, instr, substr, to_char, to_date, to_timestamp, length, lengthb, sysdate, trunc, nullif, next_day, regexp_instr, add_months, systimestamp, to_number, empty_clob and empty_blob. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● By default, the Oracle-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If

Type	Restrictions
	<p data-bbox="667 297 1422 427">newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> <li data-bbox="592 443 1366 539">● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. <li data-bbox="592 555 1414 752">● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data types supported by the primary key, see Mapping Data Types. <li data-bbox="592 768 1426 898">● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. <li data-bbox="592 913 1426 1077">● If the Oracle character set is WE8MSWIN1252 or WE8ISO8859P1, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. <li data-bbox="592 1093 1426 1189">● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. <li data-bbox="592 1205 1426 1335">● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. <li data-bbox="592 1350 1426 1447">● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. <li data-bbox="592 1462 1426 1592">● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. <li data-bbox="592 1608 1426 1872">● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. <li data-bbox="592 1888 1426 1984">● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. ● When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. ● If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the destination database sequence maximum value plus the

Type	Restrictions
	<p>security margin, and the auto-decrement sequence value is the destination database sequence minimum value minus the security margin. The default security margin is 10,000.</p> <ul style="list-style-type: none"> • If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization. • In the Oracle database, if the default value is followed by a comment starting with --, a space must be added after --. Otherwise, a syntax error will be reported when the value is migrated to the MySQL database because the MySQL database requires that -- be followed by a space.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-52 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 5-59 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-53 Synchronization instance details

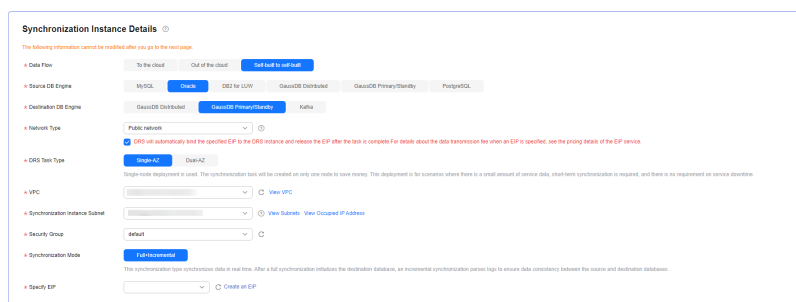


Table 5-60 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	Available options: Full , Incremental , and Full +Incremental .

Parameter	Description
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-54 Task type




Table 5-61 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p>

- Enterprise Project and Tags

Figure 5-55 Enterprise Project and Tags




Table 5-62 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-56 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name

PDB Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, related parameters have been correctly configured, and an SSL certificate has been uploaded.

Encryption Certificate

Table 5-63 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Figure 5-57 Destination database information

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-64 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-58 Synchronization mode

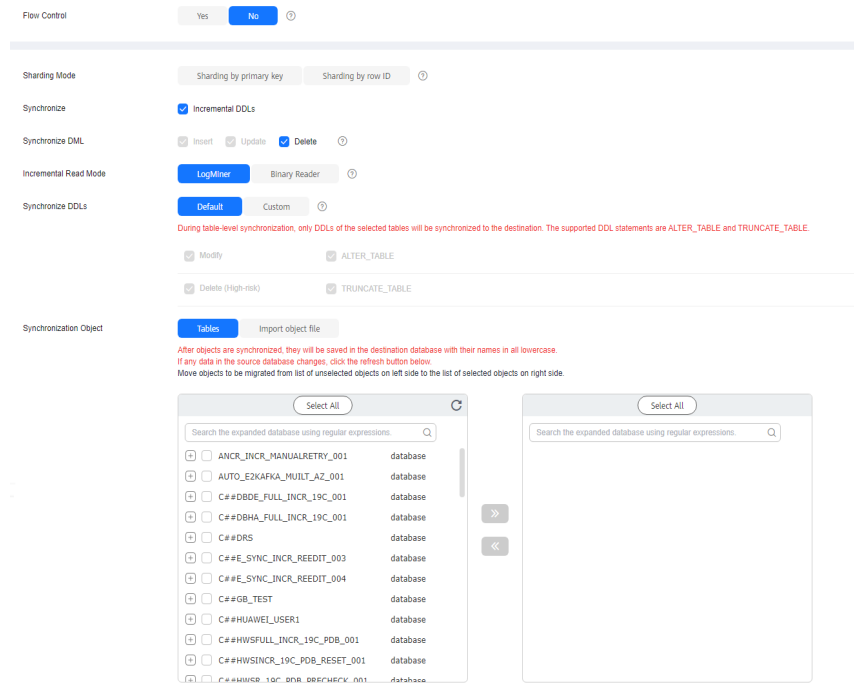
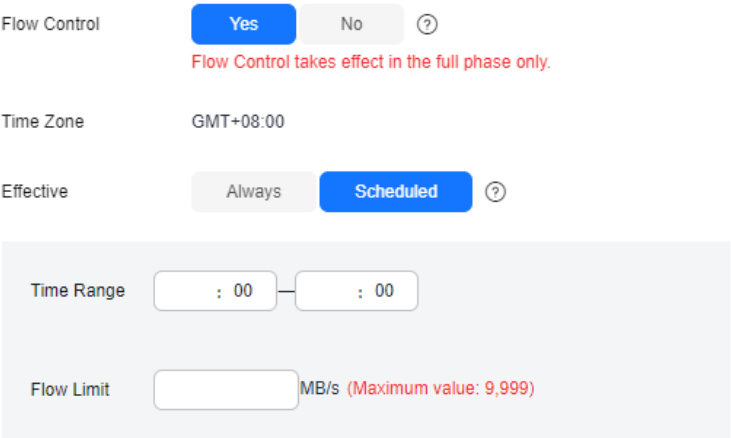



Table 5-65 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-59 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.


Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p>
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Incremental Read Mode	<p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • Binary Reader: uses DRS-developed method to directly read and parse original redo logs. The performance is high, and the Oracle resource consumption is low. <p>LogMiner is recommended.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-60 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

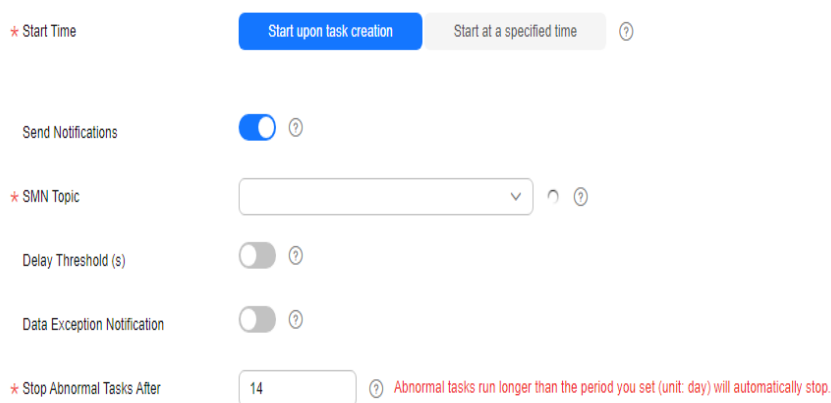
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-61 Task startup settings



* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ?

Delay Threshold (s) ?

Data Exception Notification ?


* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 5-66 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.7 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 5-67 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)	GaussDB distributed

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-68](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-68 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Source database user	<ul style="list-style-type: none"> • To synchronize a pluggable database (PDB) of Oracle 12c or later in tenant mode, the user must have the following permissions: <ul style="list-style-type: none"> - To synchronize a PDB as a PDB user, the user must have the CREATE SESSION, SELECT ANY DICTIONARY and SELECT ANY TRANSACTION permissions for the PDB and the SELECT permission for a single table. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser;) - To synchronize a PDB as a CDB user, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION permissions for the PDB, the SELECT permission for a single table, and 	<ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
	<p>CREATE SESSION, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, and SET CONTAINER permissions for the CDB.</p> <ul style="list-style-type: none"> - Authorization statement: GRANT SET CONTAINER TO <userName> CONTAINER=ALL; • In other cases: The user must have the CREATE SESSION and SELECT ANY DICTIONARY permissions. The SELECT permission on a single table is also required. Authorization statement: GRANT SELECT ON <userName.tbName> to drsUser; 	<p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p>

Type	Full Synchronization	Full+Incremental Synchronization and Incremental Synchronization
Destination database user		<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-69](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-69 Supported synchronization objects

Type	Precautions
Synchronization objects	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● Full synchronization does not support the following column types: bfile, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, sdo_geometry, urowid, interval year to month, interval day to second, and user-defined types. ● The maximum precision supported by timestamp and interval day to second is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● For the TIMESTAMP WITH TIME ZONE data type, the value cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 5-70](#).

Table 5-70 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"><li data-bbox="587 342 1385 472">• The total index length of columns in the source database cannot exceed the length limit in the destination database. For detailed length requirements, see Index Length Description.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • Do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • Do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • The rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • The username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. <p>Full synchronization</p> <ul style="list-style-type: none"> • Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. • Full synchronization consists of two phases: table structure synchronization (including indexes) and data synchronization. If the structure of a table is created in the destination database, data synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. • During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. • During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • In a full synchronization for the table structure, the length of the char and varchar2 types in the source database automatically increases by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize

Type	Restrictions
	<p>the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During table structure synchronization in a full synchronization, only default value constraints of the character string and number types are supported. Default value constraints of the function and sequence types are not supported. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● During an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. - During incremental synchronization, tables cannot be renamed, tables cannot be redefined online, and partitions cannot be swapped. ● During incremental synchronization, modify the start point of a capture task to resynchronize data. <ul style="list-style-type: none"> - After the capture point is modified, the previous object-level comparison result is not displayed. - If the start point of a capture task is changed separately, the change is synchronized to the start point of the replay task. That is, the start point of the replay task is the same as that of the capture task. This does not affect the change of the start point of the replay task. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not

Type	Restrictions
	<p>generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency.</p> <ul style="list-style-type: none">• During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999.• For incremental synchronization, the LOB type supports only the BasicFiles attribute. It does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB.• During incremental synchronization, if the source or destination database is abnormal, the task fails. After the database is recovered, the task is globally started. The status of the original capture or replay component is ignored, and the replay component is started from the capture interruption point.• In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. <p>Troubleshooting</p> <ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● The following functions can be used as default values during table structure synchronization: decode, nvl, nvl2, instr, substr, to_char, to_date, to_timestamp, length, lengthb, sysdate, trunc, nullif, next_day, regexp_instr, add_months, systimestamp, to_number, empty_clob and empty_blob. To use these functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: <ul style="list-style-type: none"> - The default value may be left empty. - The table fails to be created. As a result, the object comparison result is inconsistent or the task fails. ● If a table contains only LOB columns, data inconsistency or task failure may occur. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● By default, the Oracle-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Incremental synchronization does not support distributed transactions (XA transactions) and PARALLEL DML on an Oracle database. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. ● Ensure that a database named in lowercase letters has been created in destination database. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If

Type	Restrictions
	<p data-bbox="667 297 1422 427">newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> <li data-bbox="592 443 1366 539">● Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. <li data-bbox="592 555 1414 752">● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. For details about the data types supported by the primary key, see Mapping Data Types. <li data-bbox="592 768 1426 898">● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) storage of a single data record. <li data-bbox="592 913 1426 1077">● If the Oracle character set is WE8MSWIN1252 or WE8ISO8859P1, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. <li data-bbox="592 1093 1426 1189">● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. <li data-bbox="592 1205 1426 1335">● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a table or column name contains no more than 30 characters during an incremental synchronization. <li data-bbox="592 1350 1426 1447">● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. <li data-bbox="592 1462 1426 1592">● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. <li data-bbox="592 1608 1426 1872">● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. <li data-bbox="592 1888 1426 2013">● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the destination database is a distributed GaussDB database, the update operation performed on the source distribution column will fail to be synchronized to the GaussDB database during incremental synchronization, causing data inconsistency. Therefore, avoid updating a distribution column. ● You are advised to disable the global secondary index (GSI) function for the destination database. Otherwise, incremental synchronization may fail. ● You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. ● When editing the task to add a new table, ensure that transactions of the new table have been committed.

Type	Restrictions
	<p>Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours.</p> <ul style="list-style-type: none"> • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. • If the supplemental logging level of the source Oracle database is not set to All and the incremental update data of the source database is not found in the destination database, the complete data of the source database cannot be written to the destination database even if the conflict policy is set to Overwrite. To write all data to the destination database, set the supplemental logging level of the source database to All. • When you update a unique constraint column in an Oracle database and multiple rows are hit, data conflicts may occur. In extreme scenarios, there may be data inconsistency. • The XMLTYPE type is converted to CLOB or BLOB in the Oracle database. The usage restrictions are the same as those of CLOB and BLOB. • Only Oracle 11g and later versions support XMLTYPE data synchronization.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-62 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: ⓘ

Description: ⓘ 0/256

Table 5-71 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-63 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow: To the cloud Out of the cloud [Self-built to self-built](#)

• Source DB Engine: MySQL Oracle DB2 for LUW GaussDB Distributed GaussDB Primary Standby PostgreSQL

• Destination DB Engine: GaussDB Distributed GaussDB Primary Standby Kafka

• Network Type: Public network ⓘ

DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the DRS service.

• DRS Task Type: Single-AZ Dual-AZ Single-AZ deployment is used. This synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime.

• VPC: [Dropdown menu] [View VPC](#)

• Synchronization Instance Subnet: [Dropdown menu] ⓘ [New Subnet](#) [View Occupied IP Address](#)

• Security Group: ⓘ

• Synchronization Mode: Full Full Incremental This synchronization type synchronizes data in real time. After a full synchronization releases the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Specify EP: [Dropdown menu] [Create an EP](#)

Table 5-72 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .

Parameter	Description
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB Distributed .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	Available options: Full , Incremental , and Full +Incremental .

Parameter	Description
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-64 Task type




Table 5-73 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p>

- Enterprise Project and Tags

Figure 5-65 Enterprise Project and Tags




Table 5-74 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-66 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

PDB Name ⓘ

Database Username

Database Password

SSL Connection

Table 5-75 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

Figure 5-67 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-76 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-68 Synchronization mode

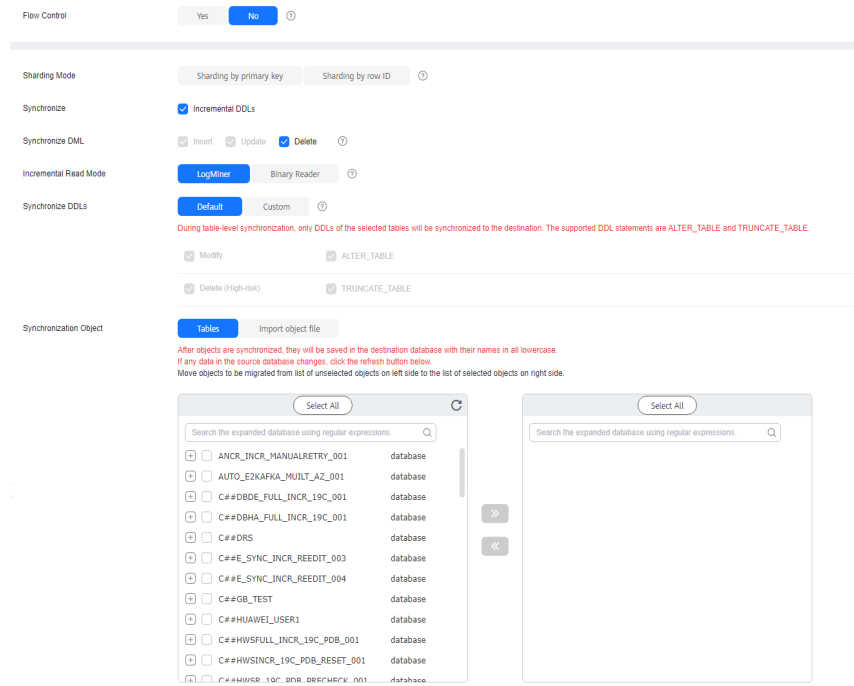
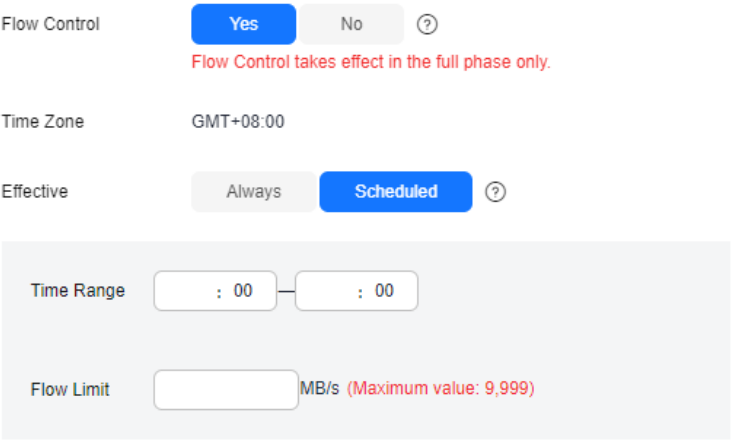



Table 5-77 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-69 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

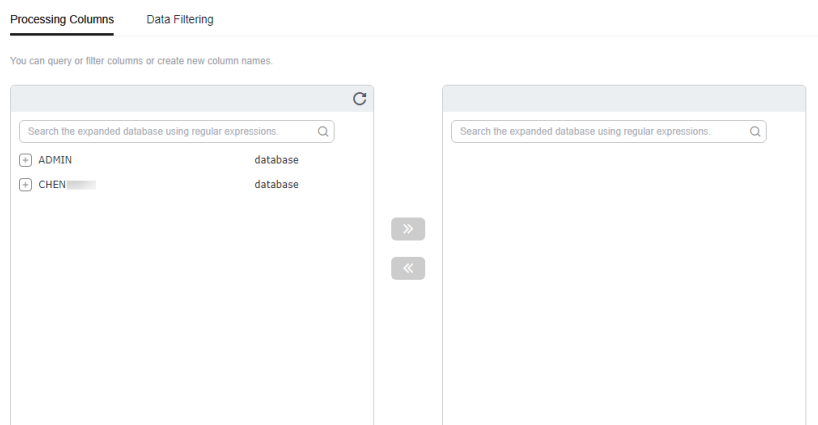
Parameter	Description
Sharding Mode	<p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p>
Synchronize	<p>Incremental DDL synchronization is supported. You can determine whether to synchronize DDLs based on service requirements.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Incremental Read Mode	<p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • Binary Reader: uses DRS-developed method to directly read and parse original redo logs. The performance is high, and the Oracle resource consumption is low. <p>LogMiner is recommended.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-70 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

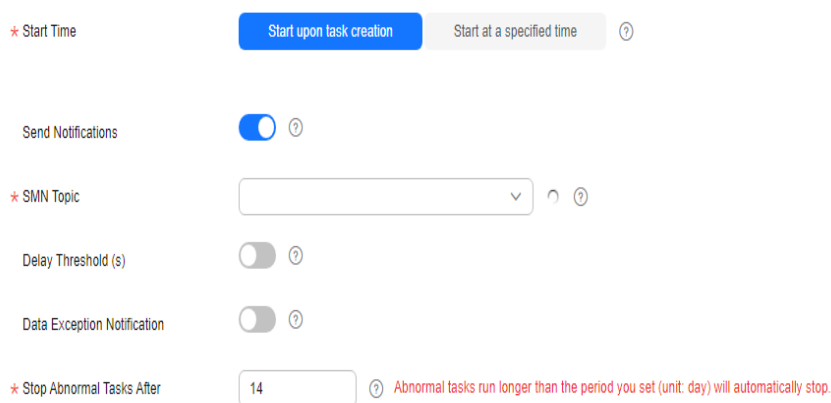
- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-71 Task startup settings



* Start Time Start upon task creation Start at a specified time ?

Send Notifications ?

* SMN Topic ? ?

Delay Threshold (s) ?

Data Exception Notification ?


* Stop Abnormal Tasks After ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 5-78 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.8 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 5-79 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) 	<ul style="list-style-type: none"> Kafka Version 0.11 or later

Supported Synchronization Objects

Table 5-80 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-80 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases. Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. <p>NOTE The restrictions on synchronization object names are as follows: The database name cannot contain +"%\<>, the schema name and table name cannot contain ".\<>, and the column name cannot contain double quotation marks (") and single quotation marks (').</p> <ul style="list-style-type: none"> Scope of incremental synchronization <ul style="list-style-type: none"> Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. Not supported: DDL statements, DML statements of unlogged tables and temporary tables <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-81](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-81 Database account permission

Type	Incremental Synchronization
Source database user	The CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, and the permission to create replication connections

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-82 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> The wal_level value of the source database must be logical. The test_decoding plug-in has been installed on the source database. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - During database-level synchronization, incremental synchronization tasks cannot be edited. - During table-level synchronization, incremental synchronization tasks can be edited, but the database cannot be changed. - Before starting a synchronization task, ensure that no long transaction is started in the source database. Otherwise, the creation of the logical replication slot will be blocked, leading the task to fail. - After a task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not delete the primary key of the source database table. Otherwise, incremental data may be lost or the task may fail. ● Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be lost or the task may fail. ● Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, incremental data may be lost or the task may fail. ● During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail. ● If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.
Stopping a task	<ul style="list-style-type: none"> ● Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. ● Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.

- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-72 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 5-83 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-73 Synchronization instance details

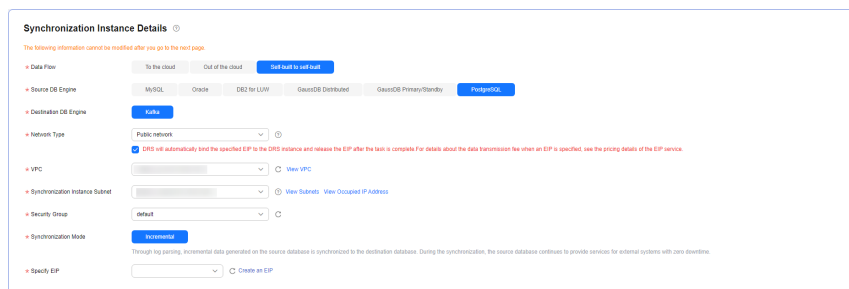


Table 5-84 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select Kafka .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-74 Task type




Table 5-85 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-75 Enterprise Project and Tags




Table 5-86 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-76 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

Table 5-87 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-77 Destination database information

Destination Database

IP Address

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol

Table 5-88 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 5-78 Synchronization mode

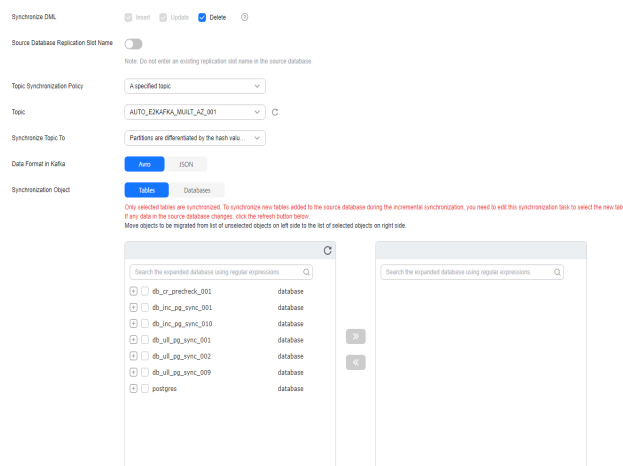



Table 5-89 Synchronization object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.

Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p>
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>

Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database, schema and table names, the performance on a single table query can be improved. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash value of the database_name.schema.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Data Format in Kafka	<p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro: A binary encoded format that is efficient. You need to deserialize the data later. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-79 Task startup settings

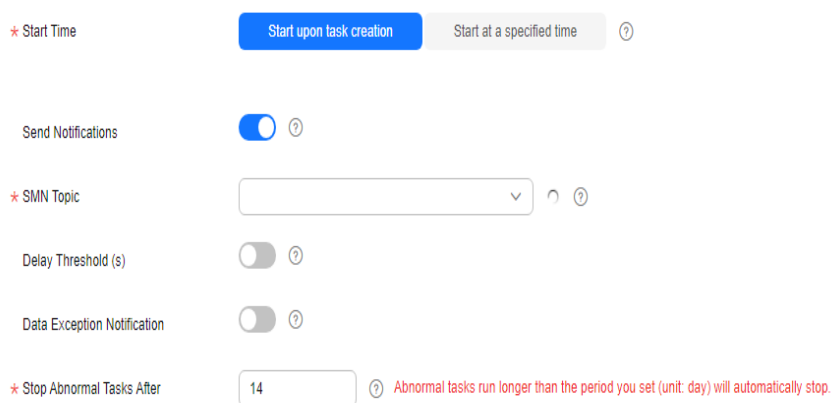



Table 5-90 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.9 From PostgreSQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-91 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14)ECS-hosted database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14)	GaussDB Primary/Standby

 NOTE

Only whitelisted users can use this function.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-92](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-92 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> 1. Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes.
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If <code>gs_loader</code> is used to create system catalogs (such as <code>public.pgxc_copy_error_log</code> and <code>public.gs_copy_summary</code>) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-93](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-93 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> - GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, real, double, bit, bit varying, boolean, bytea, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestamptz, interval, cidr, path, box, lseg, macaddr, point, polygon, inet, tsquery, tsvector, uuid, json and jsonb, are supported. - XML, line, domain, and self-built data type synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> - Supported: schemas, tables, primary key and unique constraints, table data, and sequences - Not supported: system schemas and system catalogs (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain + "%\<>", the schema name and table name cannot contain ".\<>", and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. ● Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> ● Scope of incremental synchronization

Type	Precautions
	<ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes <code>pg_logical_emit_message</code> to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-94 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The partition table trigger of the source database cannot be set to disable. - For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database. - To perform incremental synchronization: The pg_hba.conf file of the source database contains the following configuration: <pre>host replication all 0.0.0.0/0 md5</pre> The wal_level value of the source database must be logical. The test_decoding plug-in must be installed on the source database in advance. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Source database object requirements: <ul style="list-style-type: none"> - The objects that have dependencies must be synchronized at the same time. Otherwise, the synchronization may fail. <p>NOTE Common dependencies: tables referenced by views, views referenced by views, tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.</p> ● Destination database parameter requirements: <ul style="list-style-type: none"> - The max_prepared_transactions value of the destination database must be greater than that of the source database. - The max_worker_processes value of the destination database must be greater than that of the source database. - The max_locks_per_transaction value of the destination database must be greater than that of the source database.

Type	Constraints
	<ul style="list-style-type: none"> - The max_connections value of the destination database must be greater than that of the source database. - The lc_monetary value of the destination database must be the same as that of the source database. ● Destination database object requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. - The destination table can contain more columns than the source table. However, the following failures must be avoided: <p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - Only the primary table can be synchronized. The primary table will be converted into a common table and synchronized to the destination database. Data in the partition table will be written to the primary table. - When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value +original index name (which may be truncated)+_key The hash value is calculated based on the original schema

Type	Constraints
	<p>name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key.</p> <ul style="list-style-type: none"> - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. - Case conversion is not performed on names of objects after they are synchronized to the destination database.
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, data may be inconsistent or the task may fail. ● During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<ul style="list-style-type: none"> ● Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. ● Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-80 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-95 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 5-81 Synchronization instance information

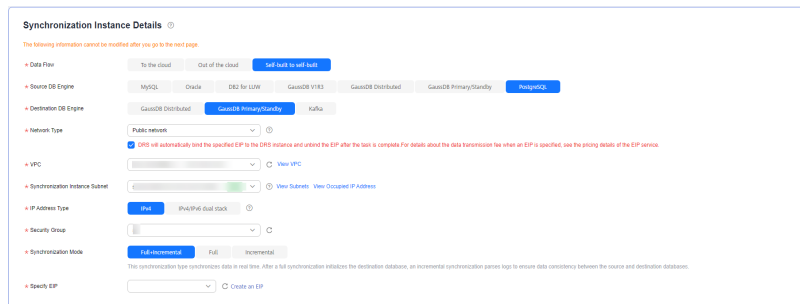


Table 5-96 Synchronization instance settings

Parameter	Description
Data Flow	Select Self-built to self-built .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	Available options: Public network and VPN or Direct Connect . Public network is used as an example.
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full This synchronization mode is suitable for scenarios where service interruption is acceptable. Database objects and data, including tables, views, and stored procedures, from non-system databases can be synchronized to the destination all at once.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-82 AZ

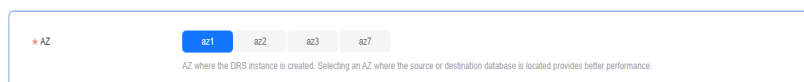


Table 5-97 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-83 Enterprise Project and Tags




Table 5-98 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-84 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the source database, and related parameters have been correctly configured.

Table 5-99 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

Figure 5-85 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

Table 5-100 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-86 Synchronization mode


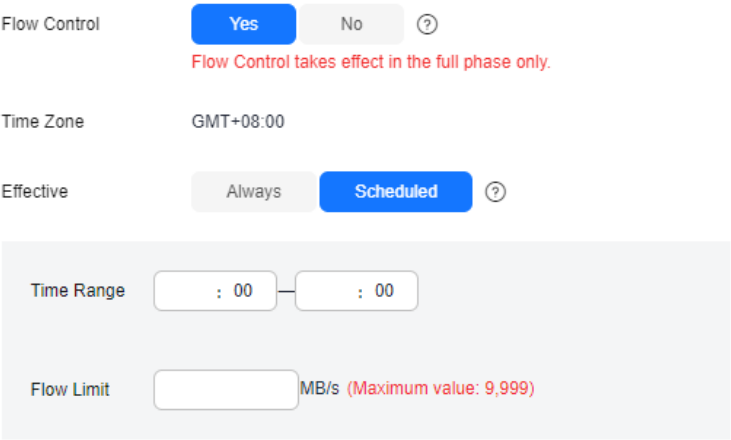



Table 5-101 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-87 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-88 Task startup settings





Table 5-102 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.10 From PostgreSQL to GaussDB Distributed

Supported Source and Destination Databases

Table 5-103 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) 	GaussDB Distributed

 **NOTE**

Only whitelisted users can use this function.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-104](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-104 Database account permissions

Type	Full Synchronization	Full+Incremental Synchronization
Source database user	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences	<p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full Synchronization	Full+Incremental Synchronization
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

[Table 5-105](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-105 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> - GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, real, double, bit, bit varying, boolean, bytea, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestamptz, interval, cidr, path, box, lseg, macaddr, point, polygon, inet, tsquery, tsvector, uuid, json and jsonb, are supported. - XML, line, domain, and self-built data type synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> - Supported: schemas, tables, primary key and unique constraints, table data, and sequences - Not supported: system schemas and system catalogs (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail.

Type	Precautions
	<p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> • Object name: The database name cannot contain +'%'\<>, the schema name and table name cannot contain ".'\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). • Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. • Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> • Scope of incremental synchronization <ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. <p>NOTE Incremental heartbeat information sending: For PostgreSQL 9.6 and later versions, if no data is written to the source database for a long time (more than 10 hours), the log extraction process invokes pg_logical_emit_message to insert heartbeat information into WAL logs of the source database, which ensures that the logical replication slot number is updated normally. (Only WAL logs are added, and services are not affected.)</p>

Precautions

The full+incremental synchronization process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-106 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The partition table trigger of the source database cannot be set to disable. - For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database. - To perform incremental synchronization: The pg_hba.conf file of the source database contains the following configuration: <pre>host replication all 0.0.0.0/0 md5</pre> The wal_level value of the source database must be logical. The test_decoding plug-in must be installed on the source database in advance. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Source database object requirements: <ul style="list-style-type: none"> - The objects that have dependencies must be synchronized at the same time. Otherwise, the synchronization may fail. <p>NOTE Common dependencies: tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.</p> ● Destination database parameter requirements: <ul style="list-style-type: none"> - The max_prepared_transactions value of the destination database must be greater than that of the source database. - The max_worker_processes value of the destination database must be greater than that of the source database. - The max_locks_per_transaction value of the destination database must be greater than that of the source database. - The max_connections value of the destination database must be greater than that of the source database.

Type	Constraints
	<ul style="list-style-type: none"> - The lc_monetary value of the destination database must be the same as that of the source database. ● Destination database object requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. - The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - Cascading one-way synchronization is not supported. For example, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - Only the primary table can be synchronized. The primary table will be converted into a common table and synchronized to the destination database. Data in the partition table will be written to the primary table. - When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value +original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_

Type	Constraints
	<ul style="list-style-type: none"> + hash value + original constraint name (which may be truncated) + _key. - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. - Case conversion is not performed on names of objects after they are synchronized to the destination database. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails.
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, data may be inconsistent or the task may fail. • During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail.
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<ul style="list-style-type: none"> ● Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. ● Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-89 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-107 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-90 Synchronization instance details

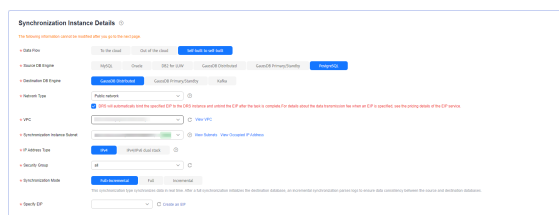


Table 5-108 Synchronization instance settings

Parameter	Description
Data Flow	Select Self-built to self-built .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB Distributed .
Network Type	Available options: Public network and VPN or Direct Connect . Public network is used as an example.
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full This synchronization mode is suitable for scenarios where service interruption is acceptable. Database objects and data, including tables, views, and stored procedures, from non-system databases can be synchronized to the destination all at once.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-91 AZ

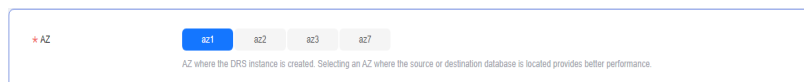


Table 5-109 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-92 Enterprise Project and Tags




Table 5-110 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-93 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

Table 5-111 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

Figure 5-94 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

Table 5-112 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-95 Synchronization mode

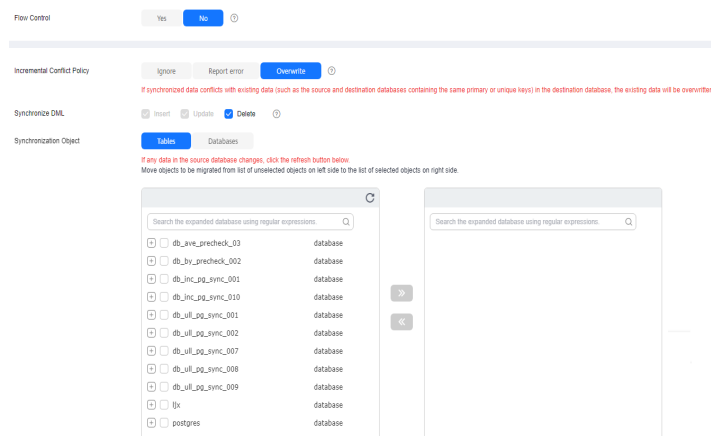
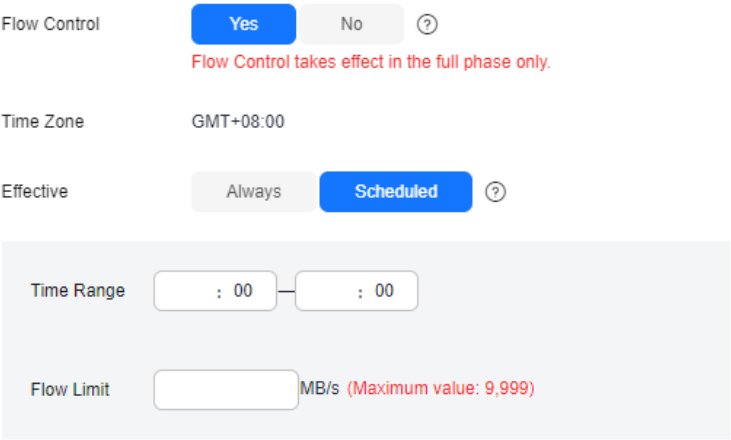



Table 5-113 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-96 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. This parameter is available for incremental tasks. After replication slot is enabled, enter the replication slot name. The name can contain a maximum of 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-97 Task startup settings





Table 5-114 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.11 From GaussDB Primary/Standby to MySQL

Supported Source and Destination Databases

Table 5-115 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<ul style="list-style-type: none"> On-premises MySQL 5.5, 5.6, and 5.7 databases ECS-hosted MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on other clouds

Supported Synchronization Objects

Table 5-116 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-116 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.)

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-117](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-117 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-118 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● The destination database parameters must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The source database supports the UTF8 character set, and the destination database supports the UTF8MB4 character set. If the source database uses the UTF8 character set and the destination database uses the UTF8MB3 character set, or if the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to inconsistent character set encoding ranges (GB18030_2022 > GB18030 > GBK > GB2312) or database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● The destination database object must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before synchronization, ensure that the table structure of the destination database has been created and is the same as that of the source database or contains all columns in the source database. - The destination table can contain more columns than the source table. However, the following failures must be avoided:

Type	Restrictions
	<p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - By default, the GaussDB-to-MySQL synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - In GaussDB, if there is 0x00 in character or binary data, the data will be truncated, which is inconsistent with MySQL. As a result, the final data may be inconsistent.

Type	Restrictions
	<ul style="list-style-type: none"> - Binary data cannot be used as a DML matching condition. That is, binary data cannot be used as a primary key. Binary data is fixed in length and is automatically filled in. The fill character of GaussDB is 0x20, and that of MySQL is 0x00. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • Replication of interval partition tables is not supported. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Do not limit the synchronization speed during data comparison. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent.

Type	Restrictions
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-98 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name:

Description:

0/256

Table 5-119 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-99 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud Out of the cloud Self-built to self-built

Source DB Engine: MySQL Oracle DB2 for LUW GaussDB Distributed GaussDB Primary/Standby PostgreSQL

Destination DB Engine: MySQL Oracle GaussDB DWS GaussDB V103 GaussDB Distributed GaussDB Primary/Standby Kafka PostgreSQL

Network Type: Public network VPC

VPC:

Synchronization Instance Subnet:

Security Group:

Synchronization Mode: Full Incremental

Specify EP:

Table 5-120 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select MySQL .

Parameter	Description
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-100 Task type




Table 5-121 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-101 Enterprise Project and Tags




Table 5-122 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-102 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

DR Synchronization

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 5-123 Source database settings

Parameter	Description
DR Synchronization	If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name .
IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-103 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password ⓘ

SSL Connection

This button is available only after the replication instance is created successfully.

Table 5-124 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-104 Synchronization mode

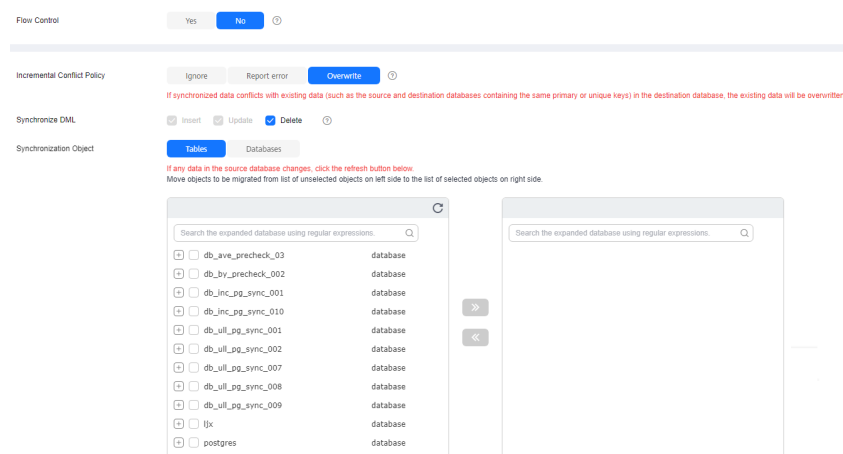
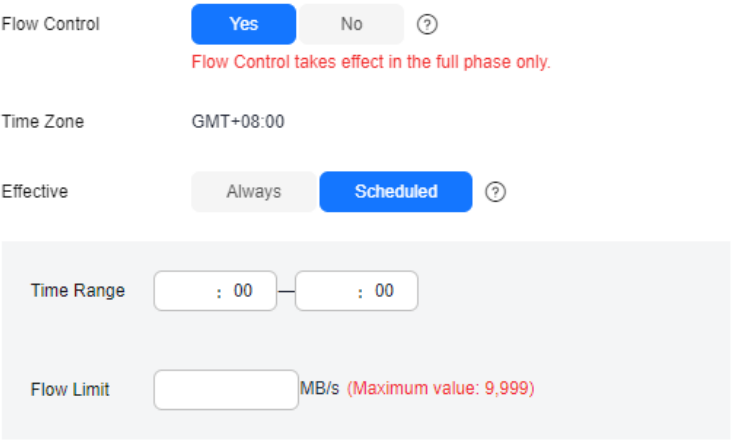



Table 5-125 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-105 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none">● Ignore The system will skip the conflicting data and continue the subsequent synchronization process.● Overwrite Conflicting data will be overwritten.● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-106 Task startup settings

Table 5-126 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.12 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 5-127 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<ul style="list-style-type: none">• On-premises Oracle 11g and 19c databases• Oracle 11g and 19c databases on an ECS

Supported Synchronization Objects

Table 5-128 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-128 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML and some DDL statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Type	Synchronization Scope
	<ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-129](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-129 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role.	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role	The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-130 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the AL32UTF8 or UTF8 character set. If the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the ZHS16GBK, ZHS16CGB231280, or ZHS32GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● Destination database object requirements: <ul style="list-style-type: none"> - Before synchronization, ensure that the corresponding table structure has been created in the destination database and is the same as that in the source instance. If column processing is performed, ensure the consistency of the table structure after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination database has sufficient disk space. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.

Type	Restrictions
	<p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - Synchronization from GaussDB primary/standby to Oracle is the backward synchronization process for synchronization from Oracle to GaussDB primary/standby. You are not advised to use either of the two synchronization modes independently. - By default, the GaussDB-to-Oracle synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The selected source database cannot contain tables with the same name but different cases. Otherwise, the synchronization fails. You are advised to synchronize only the schema and table names that are in uppercase on the Oracle database, and the schema and table names that are in lowercase on the GaussDB primary/standby database. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.

Type	Restrictions
	<ul style="list-style-type: none"> - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● For tables containing large fields (including blob, clob, nlob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. ● If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type. ● The following types of DDL statements cannot be synchronized: ALTER TABLE MERGE PARTITIONS and CREATE INDEX CONCURRENTLY. ● To synchronize DDL statements, ensure that the compatibility modes of the source and destination databases are the same. ● For DDL statements involving tablespace operations, the user of the destination database must have the tablespace operation permissions.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out.

Type	Restrictions
Synchronization comparison	<ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-107 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: DRS-9678

Description: [Text area] 0/256

Table 5-131 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-108 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

• Data Flow: **Self-built to self-built**

• Source DB Engine: MySQL, Oracle, DB2 for LUW, GaussDB Distributed, GaussDB Primary/Standby, PostgreSQL

• Destination DB Engine: **Create**, GaussDB(DWS), GaussDB Primary/Standby, Kafka

• Network Type: Public network

• VPC: [Dropdown menu] **New VPC**

• Synchronization Instance Subnet: [Dropdown menu] **New Subnet** **View Occupied IP Address**

• Security Group: default

• Synchronization Mode: **Full**, Full, Incremental

• Specify EP: [Dropdown menu] **Create an EP**

Table 5-132 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .

Parameter	Description
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select Oracle .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-109 Task type




Table 5-133 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-110 Enterprise Project and Tags




Table 5-134 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-111 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

DR Synchronization ⓘ

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 5-135 Source database settings

Parameter	Description
DR Synchronization	If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name .
IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-112 Destination database information

Destination Database

IP Address or Domain Name ⓘ
For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance.

Port

Database Service Name Service Name ⓘ

Database Username

Database Password ⓘ

SSL Connection

Table 5-136 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance.
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-113 Synchronization mode

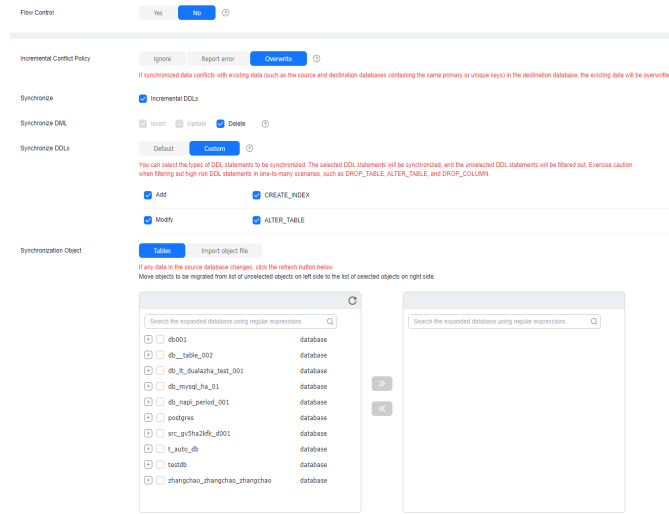
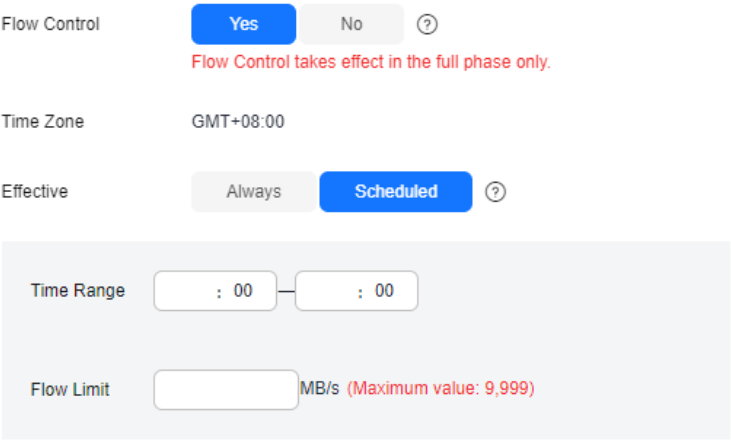



Table 5-137 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-114 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize	<p>Indicates whether to synchronize DDLs in the incremental synchronization phase.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default: DDL statements of the selected synchronization objects will be synchronized to the destination database. <ul style="list-style-type: none"> - DDL statements supported by table-level synchronization: ALTER TABLE and CREATE INDEX ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. Only selected DDL types can be synchronized.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-115 Task startup settings





Table 5-138 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.13 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 5-139 Supported databases

Source DB	Destination DB
GaussDB primary/standby	Kafka 0.11 or later

Supported Synchronization Objects

[Table 5-140](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-140 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenario: Incremental synchronization • Supported fields: BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and TINYINT • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\'\ ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-141](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-141 Database user permission

Type	Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. • The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-142 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If the DRS instance type is set to primary/standby, duplicate or inconsistent data may occur in tables without primary keys when a primary/standby switchover is performed. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: select working_version_num();

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During migration of table-level objects, you are not advised to rename the tables. • Replication of interval partition tables is not supported. • The name of a primary key column cannot be changed. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-116 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page.
 The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256

Table 5-143 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-117 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow
To the cloud
Out of the cloud
Task load is not full

Source DB Engine
MySQL
Oracle
DB2 for LUW
GaussDB Distributed
Supports Primary/Standby
PostgreSQL

Destination DB Engine
Click
GaussDB(DWS)
GaussDB Primary/Standby
Taurus

Network Type
Public network

DRS will automatically bind the specified EP to the DRS instance and release the EP after the task is completed. For details about the data transmission fee when an EP is specified, see the pricing details of the EP service.

CRS Test Type
Single AZ
Dual AZ

Single-az deployment is used. The synchronization task will be created on only one az to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service availability.

VPC
..
View VPC

Synchronization Instance Subnet
..
View Subnets
View Occupied IP Address

Security Group
..
View Security Groups

Synchronization Mode
Incremental

Through log parsing, incremental data generated on the source database is synchronized to the destination database. During the synchronization, the source database continues to provide services for external systems with zero downtime.

Security EP
..
Create an EP

Table 5-144 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select Kafka .
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
DRS Task Type	Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">– Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.– Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task .
VPC	Select an available VPC.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-118 Task type





Table 5-145 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-119 AZ</p> 

- Enterprise Project and Tags

Figure 5-120 Enterprise Project and Tags




Table 5-146 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-121 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

DR Synchronization ⓘ

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 5-147 Source database settings

Parameter	Description
DR Synchronization	If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name .
IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-122 Destination database information

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 5-148 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 5-123 Synchronization mode

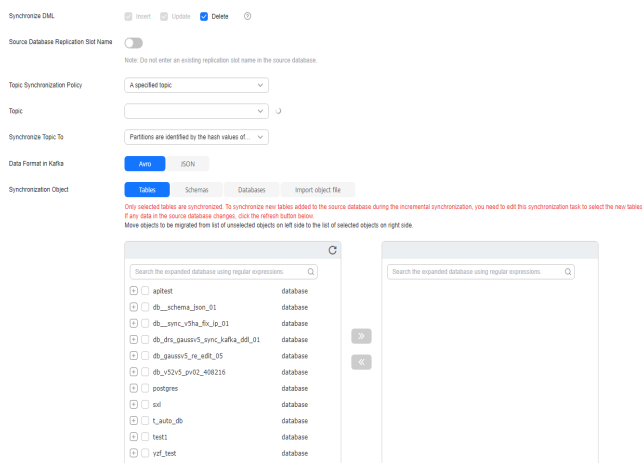



Table 5-149 Synchronization object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy is set to A specified topic .
Topic Name Format	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal. If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.

Parameter	Description
Number of Partitions	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p>
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • Partitions are differentiated by the hash values of database_name.schema_name.table_name: This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved. • Partitions are identified by the hash values of the primary key: This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.schema_name.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Data Format in Kafka	<p>Select the format of data delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro: A binary encoded format that is efficient. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-124 Task startup settings





Table 5-150 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.14 From GaussDB Primary/Standby to GaussDB Distributed

Supported Source and Destination Databases

Table 5-151 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<p>GaussDB distributed</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p>

Supported Synchronization Objects

Table 5-152 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-152 Supported synchronization objects

Type	Constraints
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none"> - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete. - The database name, schema name, and table name cannot contain special characters /<.>\\'\ ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-153](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-153 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-154 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICA IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the

Type	Constraints
	<p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance.

Type	Constraints
	<ul style="list-style-type: none"> - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not compare data during full synchronization. After full synchronization is complete, the comparison is automatically triggered. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-125 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-155 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-126 Synchronization instance details

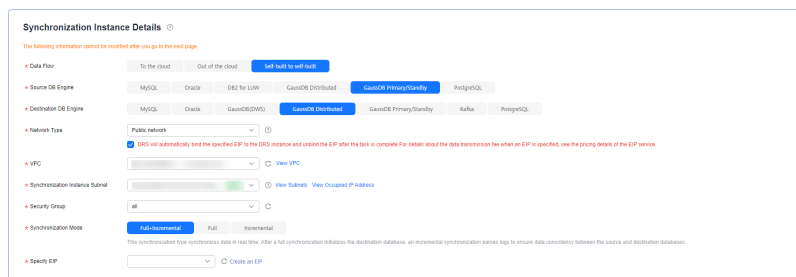


Table 5-156 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB Distributed .
Network Type	Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. If you select VPN or Direct Connect for Network Type , you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-127 Task type




Table 5-157 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-128 Enterprise Project and Tags




Table 5-158 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-129 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

DR Synchronization

IP Address or Domain Name

Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 5-159 Source database settings


Parameter	Description
DR Synchronization	If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name .
IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-130 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-160 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-131 Synchronization mode

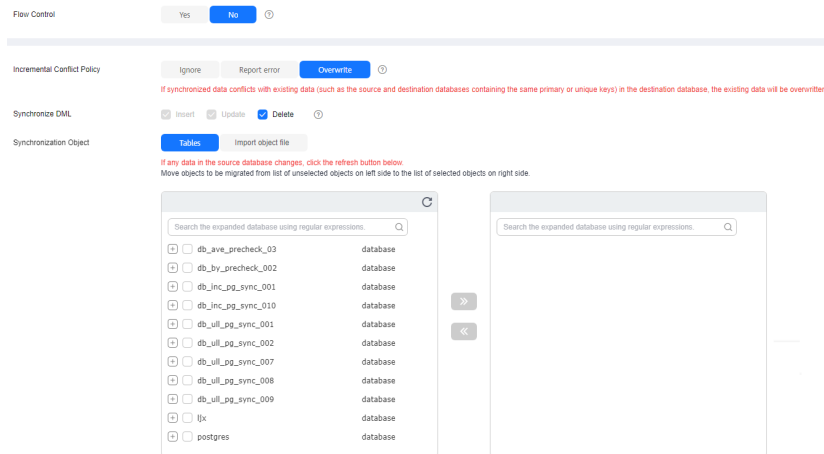
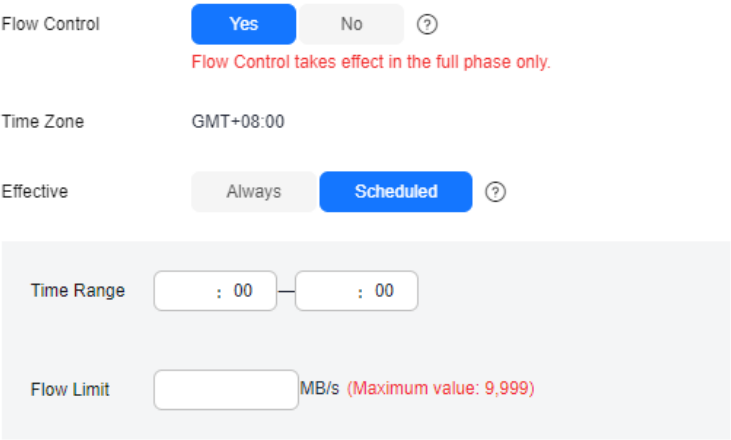



Table 5-161 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-132 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-133 Task startup settings

Table 5-162 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.15 From GaussDB Primary/Standby to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-163 Supported databases

Source DB	Destination DB
GaussDB primary/standby	GaussDB primary/standby NOTE The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

[Table 5-164](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-164 Supported synchronization objects

Type	Constraints
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none">- Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.- The database name, schema name, and table name cannot contain special characters /<.>\\' ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-165](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-165 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-166 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Select full or full+incremental synchronization. If a table structure already exists in the destination database, ensure that the table structure is the same as that on the source database. If column processing is performed, ensure that the table structure is the same as that after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the

Type	Constraints
	<p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not compare data during full synchronization. After full synchronization is complete, the comparison is automatically triggered. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-134 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 5-167 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-135 Synchronization instance details



Table 5-168 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. If you select VPN or Direct Connect for Network Type , you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-136 Task type




Table 5-169 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-137 Enterprise Project and Tags




Table 5-170 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-138 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

DR Synchronization ⓘ

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 5-171 Source database settings


Parameter	Description
DR Synchronization	If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name .
IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.


Figure 5-139 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password 

Test Connection This button is available only after the replication instance is created successfully.

Table 5-172 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-140 Synchronization mode

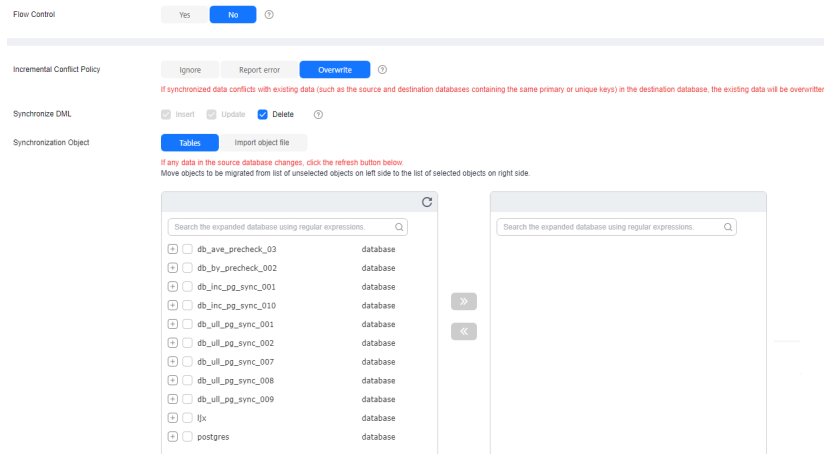
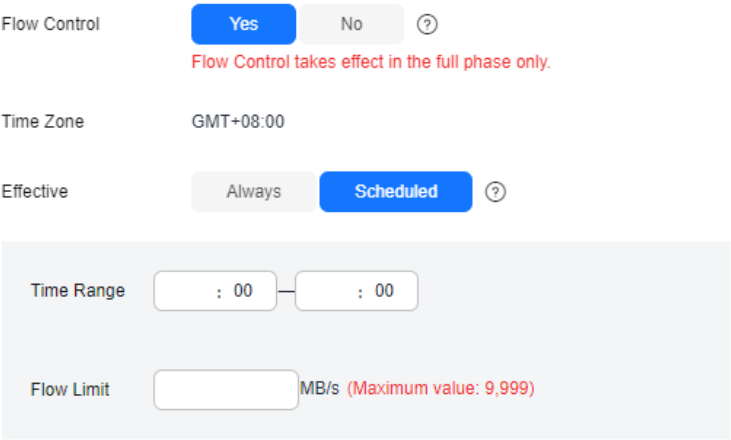



Table 5-173 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-141 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-142 Task startup settings

Table 5-174 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.16 From GaussDB Primary/Standby to PostgreSQL

Supported Source and Destination Databases

Table 5-175 Supported databases

Source DB	Destination DB
GaussDB Primary/Standby	<ul style="list-style-type: none">On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)

 NOTE

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 5-176](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-176 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● The following field types are not supported: SMALLDATETIME, REETIME, ABSTIME, TID, XID, CID and OID ● Table-level synchronization is supported. <ul style="list-style-type: none"> - Table data and sequence values can be synchronized. - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\` \?! The column name cannot contain double quotation marks (""), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 5-177](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-177 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	The user must have the following minimum permissions: <ul style="list-style-type: none"> ● Database permission: CONNECT ● Schema permission: USAGE ● Table permission: INSERT, UPDATE, DELETE, and SELECT ● Sequence permission: UPDATE 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-178 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The destination database, schemas, and table object structures must be created in advance. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Constraints
	<ul style="list-style-type: none"> - Before full synchronization, ensure that there is no data in the destination database table. Otherwise, data may be inconsistent or the task may fail. - If retry upon failure or resumable transfer occurs during full synchronization, the truncate command is executed to clear data in tables without primary keys in the destination database that have not been synchronized, and then the tables are synchronized again. Before full synchronization, ensure that the synchronization user has the TRUNCATE permission on the tables without primary keys in the destination database. Otherwise, the task may fail. ● Other notes: <ul style="list-style-type: none"> - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● When GaussDB is not compatible with PostgreSQL, the processing precision of the date data type in the GaussDB database may be different from that in the PostgreSQL database. As a result, data of the two databases is inconsistent.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-143 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name: ⓘ

Description: ⓘ

0/256

Table 5-179 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-144 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

* Synchronization Type: **One-way** Two-way

* Data Flow: **Self-built to self-built** To the cloud Out of the cloud

* Source DB Engine: MySQL Oracle **GaussDB Primary/Standby** PostgreSQL Microsoft SQL Server

* Destination DB Engine: MySQL Oracle GaussDB VRS GaussDB Distributed GaussDB Primary/Standby Kafka **PostgreSQL**

* Network Type: VPC ⓘ

* VPC: vpc-c6b-6e61-1b-3-3-085 ⓘ View VPC

* Synchronization Instance Subnet: subnet-a752-1b-83-35-8246 ⓘ The IP address is allocated automatically but it can't be changed. ⓘ View Subnets View Occupied IP Address

* IP Address Type: **Private** Public ⓘ

* Security Group: default ⓘ

* Synchronization Mode: **Full-Incremental** Full Incremental

The synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

Table 5-180 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Primary/Standby .

Parameter	Description
Destination DB Engine	Select PostgreSQL .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-145 AZ



Table 5-181 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-146 Enterprise Project and Tags




Table 5-182 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-147 Source database information

Source Database

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name: [?](#)

Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username:

Database Password:

[Test Connection](#) This button is available only after the replication instance is created successfully.

Table 5-183 Source database settings

Parameter	Description
IP Address or Domain Name	Enter the IP addresses of the GaussDB database in the IP address/Domain name:Port format. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-148 Destination database information

Destination Database

Select Connection

VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the destination database, and related parameters have been correctly configured.

This button is available only after the replication instance is created successfully.

Table 5-184 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-149 Synchronization mode

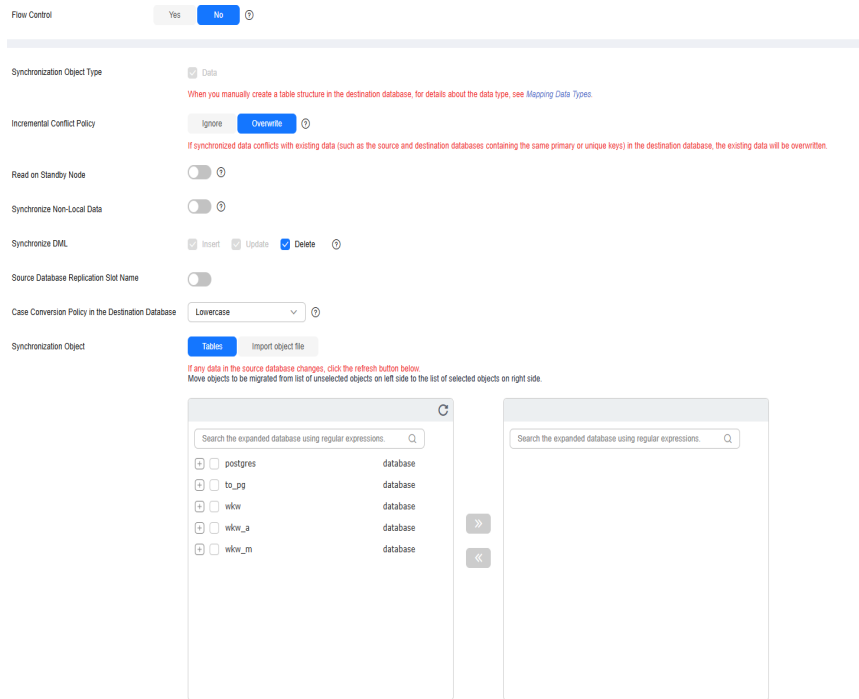
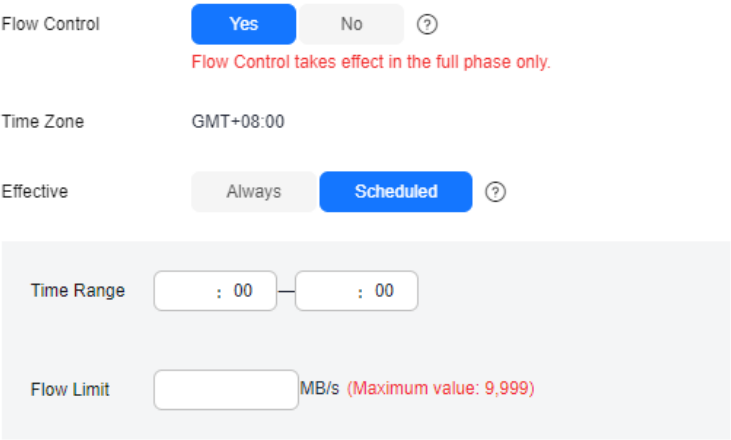



Table 5-185 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-150 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	Type of objects for full synchronization. Data is mandatory.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>
Case Conversion Policy in the Destination Database	<p>Case of schema names, table names, and column names can be converted. If you have specified the mapping name when selecting objects to be synchronized, ignore this parameter.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-151 Task startup settings





Table 5-186 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.17 From GaussDB Distributed to MySQL

Supported Source and Destination Databases

Table 5-187 Supported databases

Source DB	Destination DB
GaussDB distributed	<ul style="list-style-type: none"> On-premises MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on an ECS MySQL 5.5, 5.6, and 5.7 databases on other clouds

Supported Synchronization Objects

Table 5-188 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-188 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-189](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-189 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user must have the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		details, see Connection and Port Description for Incremental Synchronization from GaussDB.	a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES		

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
 - To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-190 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the UTF8MB4 character set. If the source database uses the UTF8 character set and the destination database uses the UTF8MB3 character set, or if the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to inconsistent character set encoding ranges (GB18030_2022 > GB18030 > GBK > GB2312) or database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● The destination database object must meet the following requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before synchronization, ensure that the table structure of the destination database has been created and is the same as that of the source database or contains all columns in the source database. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Restrictions
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - By default, the GaussDB-to-MySQL synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - The names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: select working_version_num();

Type	Restrictions
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.

Type	Restrictions
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Do not limit the synchronization speed during data comparison. MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-152 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 5-191 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-153 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

- Data Flow: To the cloud Out of the cloud Self-built to self-built
- Source DB Engine: MySQL Oracle DB2 for LUW GaussDB Distributed GaussDB Primary/Standby PostgreSQL
- Destination DB Engine: MySQL Oracle GaussDB(DWH) GaussDB Distributed GaussDB Primary/Standby Kafka PostgreSQL
- Network Type: Public network
- VPC: DRS will automatically bind the specified VPC to the DRS instance and unbind the VPC after the task is complete. For details about the data transmission fee when an VPC is specified, see the pricing details of the DRS service.
- Synchronization Instance Subnet: View Subnets
- IP Address Type: IPv4 IPv4/IPv6 dual stack
- Security Group: View VPC
- Synchronization Mode: Full Incremental
- This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization passes logs to ensure data consistency between the source and destination databases.
- Source DN Quantity:
- Specify EIP: Create an EIP

Table 5-192 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select MySQL .
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Incremental, and Full</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-154 Task type




Table 5-193 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-155 Enterprise Project and Tags




Table 5-194 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> <li data-bbox="691 297 1422 394">– Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. <li data-bbox="691 409 1406 573">– If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. <li data-bbox="691 589 1417 651">– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-156 Source database information

Source Database

CN IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ?

IP Address or Domain Name

Test Connection

Table 5-195 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-157 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 5-196 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-158 Synchronization mode

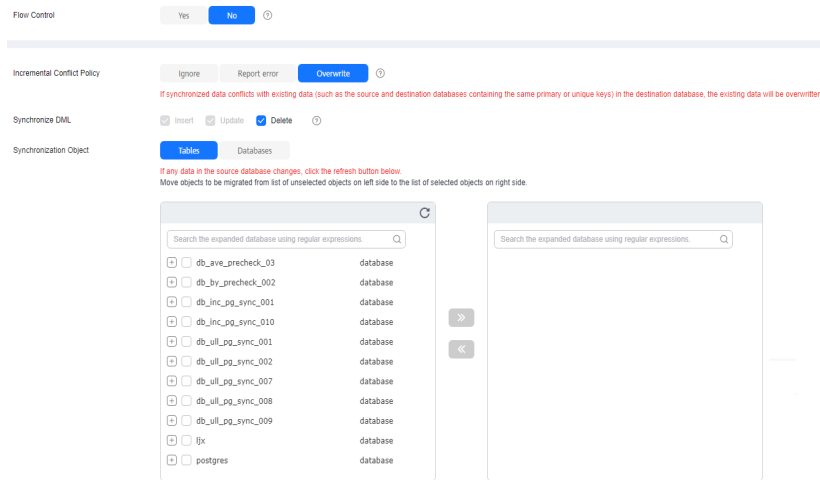
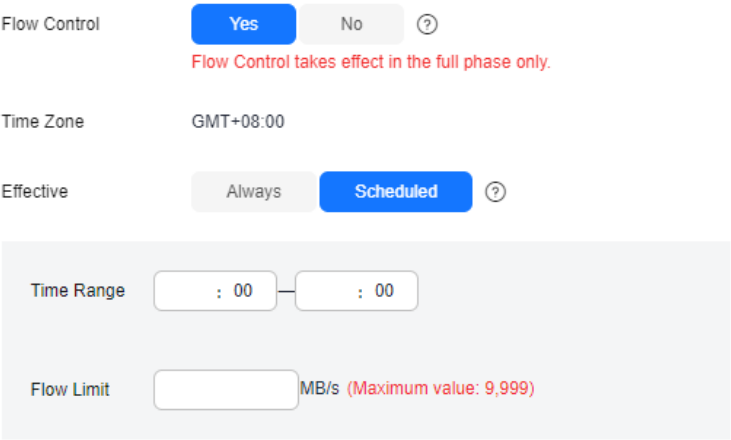



Table 5-197 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-159 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-160 Task startup settings

Table 5-198 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.18 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 5-199 Supported databases

Source DB	Destination DB
GaussDB distributed	<ul style="list-style-type: none">• On-premises Oracle 11g and 19c databases• Oracle 11g and 19c databases on an ECS

Supported Synchronization Objects

Table 5-200 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-200 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and MONEY • Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\` \?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-201](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-201 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to

Type	Full	Incremental	Full+Incremental
		<p>Connection and Port Description for Incremental Synchronization from GaussDB.</p>	<p>a GaussDB Database.</p> <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.
Destination database user	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role.</p>	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</p>	<p>The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</p>

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-202 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the AL32UTF8 or UTF8 character set. If the source database uses the GBK, GB18030, or ZHS16GBK character set and the destination database uses the ZHS16GBK, ZHS16CGB231280, or ZHS32GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data may be inconsistent after synchronization, or data may fail to be written into the destination database. ● Destination database object requirements: <ul style="list-style-type: none"> - Before synchronization, ensure that the corresponding table structure has been created in the destination database and is the same as that in the source instance. If column processing is performed, ensure the consistency of the table structure after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination database has sufficient disk space. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.

Type	Restrictions
	<p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - Synchronization from GaussDB distributed to Oracle is the backward synchronization process for synchronization from Oracle to GaussDB distributed. You are not advised to use either of the two synchronization modes independently. - By default, the GaussDB-to-Oracle synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - The selected source database cannot contain tables with the same name but different cases. Otherwise, the synchronization fails. You are advised to synchronize only the schema and table names that are in uppercase on the Oracle database, and the schema and table names that are in lowercase on the GaussDB distributed database. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in

Type	Restrictions
	<p>a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails.</p> <ul style="list-style-type: none"> - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <code>select working_version_num();</code>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● For tables containing large fields (including blob, clob, nclob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. ● If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison.

Type	Restrictions
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-161 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:
Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 5-203 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-162 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow: To the cloud Out of the cloud **Self-built to self-built**

Source DB Engine: MySQL Oracle DB2 for LUW **GaussDB Distributed** GaussDB Primary/Standby PostgreSQL

Destination DB Engine: **Oracle** GaussDB(DWS) GaussDB Distributed Kafka

Network Type: ⓘ

VPC: ⓘ

Synchronization Instance Subnet: ⓘ

Security Group: ⓘ

Synchronization Mode: **Full** Full Incremental

Source DN Quantity: ⓘ

Specify EP: ⓘ

Table 5-204 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Oracle .

Parameter	Description
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.

Parameter	Description
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- Task Type

Figure 5-163 Task type




Table 5-205 Task type information

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications .
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-164 Enterprise Project and Tags




Table 5-206 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-165 Source database information

Source Database

CN IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ?

IP Address or Domain Name

Test Connection

Table 5-207 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-166 Destination database information

Destination Database

IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance. ⓘ

Port

Database Service Name Service Name ⓘ

Database Username

Database Password

SSL Connection

Table 5-208 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance.
Port	The port of the destination database. Range: 1 - 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-167 Synchronization mode

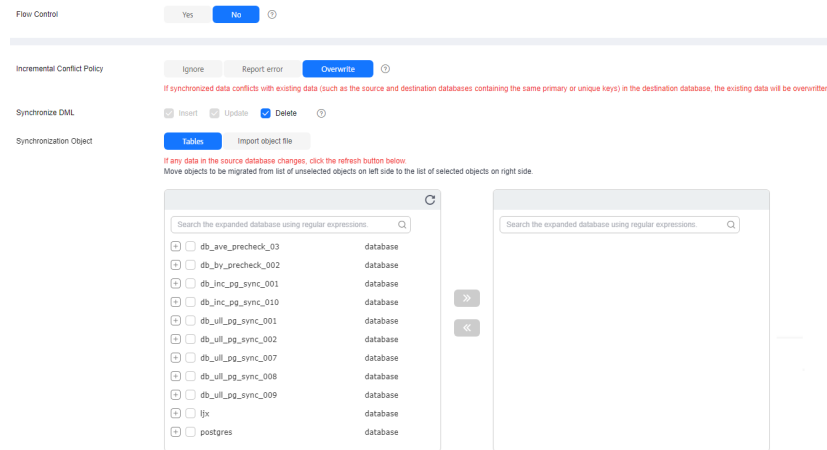
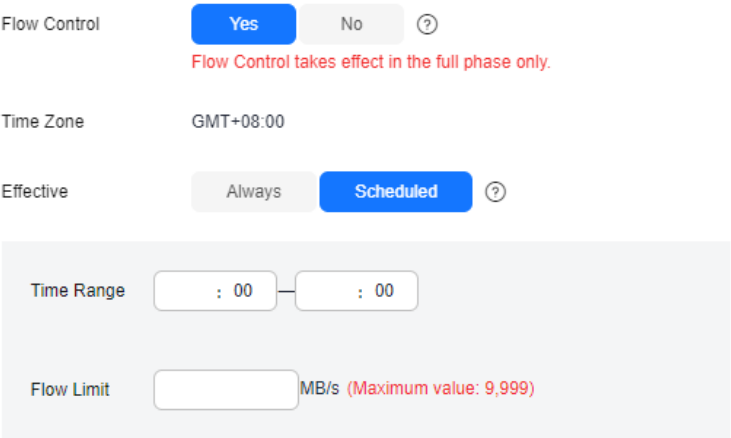



Table 5-209 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-168 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-169 Task startup settings

Table 5-210 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.19 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 5-211 Supported databases

Source DB	Destination DB
GaussDB distributed	Kafka 0.11 or later

Supported Synchronization Objects

Table 5-212 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-212 Supported synchronization objects

Type	Synchronization Scope
Synchronization scope	<ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenario: Incremental synchronization • Supported fields: BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, and TINYINT • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of the selected table can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.

Type	Synchronization Scope
	<ul style="list-style-type: none"> The database name, schema name, and table name cannot contain special characters /<.>\\'\ ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-213](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-213 Database user permission

Type	Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a

synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-214 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` \?! ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME

Type	Constraints
	<p>ZONE data cannot be synchronized. To query the source database version, run the following SQL statement:</p> <pre>select working_version_num();</pre>
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the REPLICATION IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During migration of table-level objects, you are not advised to rename the tables. • Replication of interval partition tables is not supported. • The name of a primary key column cannot be changed. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-170 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Region dropdown] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Project dropdown]

* Task Name: [Task Name input: DRS-5678] ⓘ

Description: [Description text area: 0/256] ⓘ

Table 5-215 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-171 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow: [To the cloud] | [Out of the cloud] | **Self-built to self-built**

Source DB Engine: [Kafka] | [GaussDB Distributed] | [GaussDB Primary/Standby] | [PostgreSQL]

Destination DB Engine: [Oracle] | **GaussDB(DWS)** | [GaussDB Distributed] | [Other]

Network Type: [Public network] ⓘ

DRS Task Type: [Replicate] | [Backup] ⓘ

VPC: [vpc-xxxxxxxxxxxx] ⓘ [View VPC](#)

Synchronization Instance Subnet: [subnet-xxxxxxxxxx] ⓘ [View Subnet](#)

Security Group: [sg-xxxxxxxxxx] ⓘ

Elasticsearch Instance: **Activate**

Source DB Quota: [2] ⓘ

Specify EP: [] ⓘ [Create an EP](#)

Table 5-216 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Kafka .

Parameter	Description
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-172 Task type




Table 5-217 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-173 Enterprise Project and Tags




Table 5-218 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-174 Source database information

Source Database

CN IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ?

IP Address or Domain Name

Test Connection

Table 5-219 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-175 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 5-220 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Connection Method	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 5-176 Synchronization mode

Synchronize DDL Insert Update Delete ⓘ

Source Database Replication Slot Name
Note: Do not enter an existing replication slot name in the source database.

Topic Synchronization Policy

Topic

Synchronize Topic To

Date Format in Kafka

Synchronization Object

Only selected tables are synchronized. To synchronize new tables added to the source database during the incremental synchronization, you need to edit this synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.


Search the expanded database using regular expressions.

- apptest database
- db__schema_bson_01 database
- db__sync_v5ha_fm_ip_01 database
- db_drs_gaussv5_sync_kafka_ddl_01 database
- db_gaussv5_m_461_05 database
- db_v52v5_pvt02_408216 database
- postgres database
- test database
- test1 database
- test2 database
- test3 database

Search the expanded database using regular expressions.

Table 5-221 Synchronization Object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Source Database Replication Slot Name	You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal. If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.
Number of Partitions	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.
Replication Factor	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.

Parameter	Description
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database, schema and table names, the performance on a single table query can be improved. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash value of the database_name.schema.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Data Format in Kafka	<p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

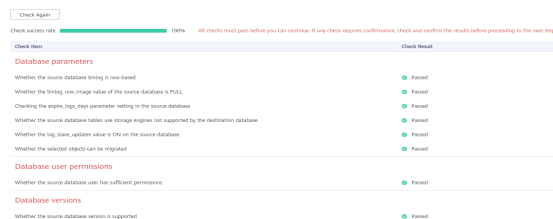
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-177 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-178 Task startup settings

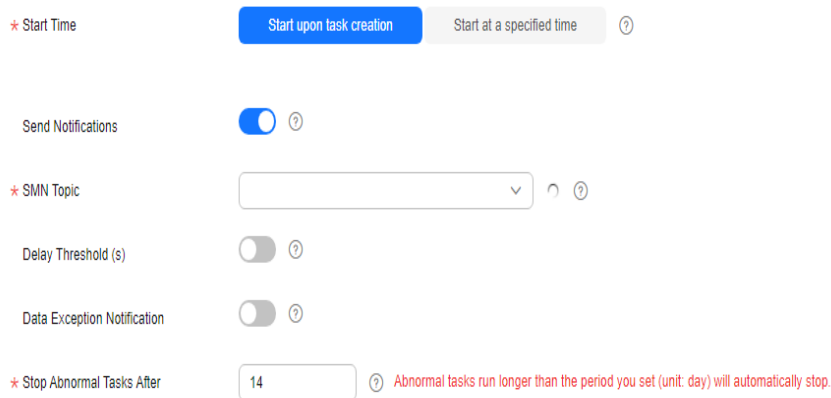



Table 5-222 Task startup settings

Parameter	Description
Start Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p>
Send Notifications	<p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.</p>
SMN Topic	<p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p>
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.20 From GaussDB Distributed to GaussDB Distributed

Supported Source and Destination Databases

Table 5-223 Supported databases

Source DB	Destination DB
GaussDB distributed	GaussDB distributed NOTE The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

[Table 5-224](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-224 Supported synchronization objects

Type	Constraints
Synchronizati on scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none">- Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.- The database name, schema name, and table name cannot contain special characters /<.>\\' ! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-225](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-225 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB .	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • To use value comparison, you must have the MONADMIN permission. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-226 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<>\\'\"?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Constraints
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - GaussDB distributed supports only ranged partitioned tables. Therefore, other types of partitioned tables are synchronized to the destination database as common tables. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. The topology structure of the source database must be the same as that of the destination database. Otherwise, the content comparison function is unavailable. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-179 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.

The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region ⌵

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project ⌵

Task Name ⓘ

DRS-5678

Description ⓘ

0/256

Table 5-227 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-180 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

Data Flow ⓘ

To the cloud
Out of the cloud
Inflight to on-prem

Source DB Engine ⓘ

MySQL
Oracle
DB2 for LUW
GaussDB (Standard)
GaussDB (Primary/Standby)
PostgreSQL

Destination DB Engine ⓘ

Oracle
GaussDB (Standard)
GaussDB (Primary/Standby)
Kafka

Network Type ⓘ

Public network
⌵

⚠ OSS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

VPC ⓘ

xxxxxxxxxxxxxxxxxxxx
⌵
New VPC

Synchronization Instance Subnet ⓘ

xxxxxxxxxxxxxxxxxxxx
⌵
New Subnet

Security Group ⓘ

default
⌵

Synchronization Mode ⓘ

Full Incremental
Full
Incremental

The synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization pattern helps to ensure data consistency between the source and destination databases.

Source DB Quota ⓘ

2
⌵

Specify EIP ⓘ

⌵
Create an EIP

Table 5-228 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB Distributed .
Network Type	Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. If you select VPN or Direct Connect for Network Type , you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-181 Task type




Table 5-229 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-182 Enterprise Project and Tags




Table 5-230 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> <li data-bbox="691 297 1420 394">– Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. <li data-bbox="691 409 1406 573">– If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. <li data-bbox="691 589 1417 649">– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-183 Source database information

Source Database

CN IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ?

IP Address or Domain Name

Test Connection

Table 5-231 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-184 Destination database information

Destination Database

IP Address or Domain Name ⓘ

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

This button is available only after the replication instance is created successfully.

Table 5-232 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-185 Synchronization mode

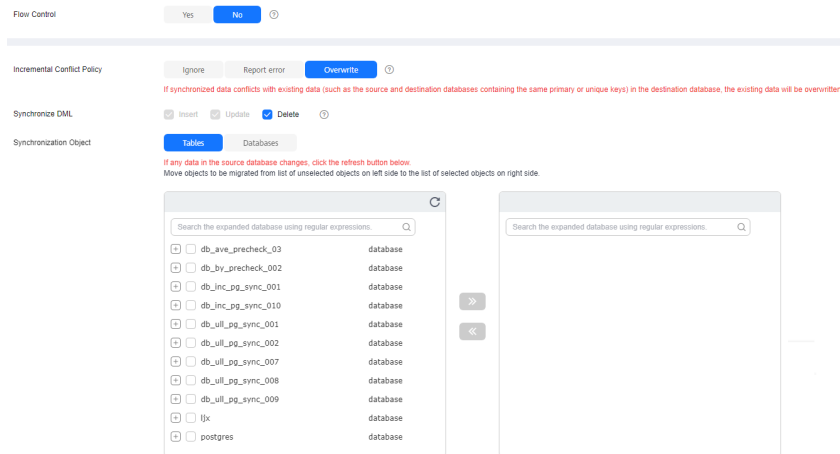
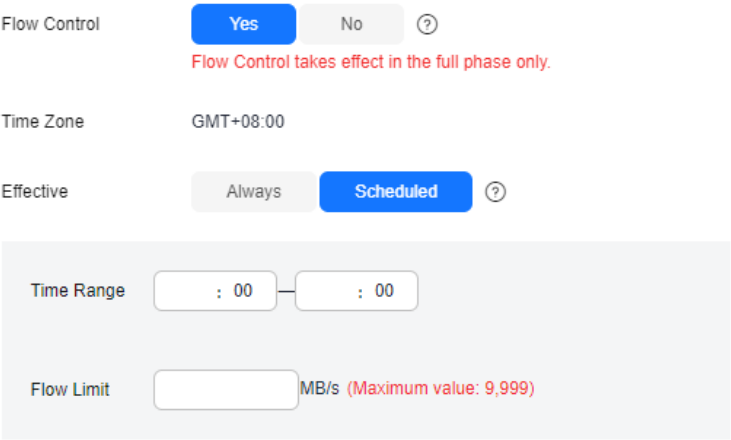



Table 5-233 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-186 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-187 Task startup settings

Table 5-234 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.21 From GaussDB Distributed to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-235 Supported databases

Source DB	Destination DB
GaussDB distributed	GaussDB primary/standby NOTE <ul style="list-style-type: none">The destination database version must be the same as or later than the source database version.

Supported Synchronization Objects

[Table 5-236](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-236 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, TINYINT UNSIGNED, SMALLINT UNSIGNED, INTEGER UNSIGNED, BIGINT UNSIGNED, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID, and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the table structure, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, replication tables without primary keys, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized.

Type	Constraints
	<ul style="list-style-type: none">- Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete.- The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user.

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-237](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-237 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-238 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records

Type	Constraints
	<p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - By default, the GaussDB-to-GaussDB synchronization task does not support loopback and cascading synchronization. That is, data cannot be synchronized from instance A to instance B and then from instance B to instance C. - During real-time synchronization, the consistency of distributed transactions is not ensured. - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - After a table is synchronized and renamed, the new index name format is <i>i_+hash value+original index name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - After a table is synchronized and renamed, the new constraint name format is <i>c_+hash value+original constraint name (which may be truncated)+_key</i>. The hash value is calculated based on the <i>original schema name_original table name_mapped schema name_mapped table name_original index name</i> - GaussDB distributed supports only ranged partitioned tables. Therefore, other types of partitioned tables are synchronized to the destination database as common tables. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to

Type	Constraints
	<p>synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> - If the B compatibility mode is enabled for the destination database, synchronizing data of the interval type may cause task failures. You are not advised to synchronize data of the interval type. - In many-to-one scenarios, ensure that the unique keys in the source database tables do not conflict. You are advised to set filtering criteria for row comparison and value comparison. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Data processing	<ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered.

Type	Constraints
Synchronization comparison	<ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison.
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-188 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
 The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

02/56

Table 5-239 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-189 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

• Data Flow To the cloud Out of the cloud [Self built to self built](#)

• Source DB Engine MySQL Oracle DB2 for LUW [GaussDB Distributed](#) GaussDB Primary/Standby PostgreSQL

• Destination DB Engine MySQL Oracle GaussDB(DWS) GaussDB Distributed [GaussDB Primary/Standby](#) Kafka PostgreSQL

• Network Type Public network ⓘ

☑ DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the DRS service.

• VPC ⓘ [View VPC](#)

• Synchronization Instance Subnet ⓘ [View Subnet](#)

• Security Group ⓘ

• Synchronization Mode [Full-replacement](#) Full Incremental

This synchronization type synchronizes data in real time. After a full synchronization initiates the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

• Source DB Quota ⓘ

• Specify EIP ⓘ [Create an EIP](#)

Table 5-240 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 5-190 Task type




Table 5-241 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-191 Enterprise Project and Tags




Table 5-242 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-192 Source database information

Source Database

CN IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ?

IP Address or Domain Name

Test Connection

Table 5-243 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-193 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-244 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-194 Synchronization mode

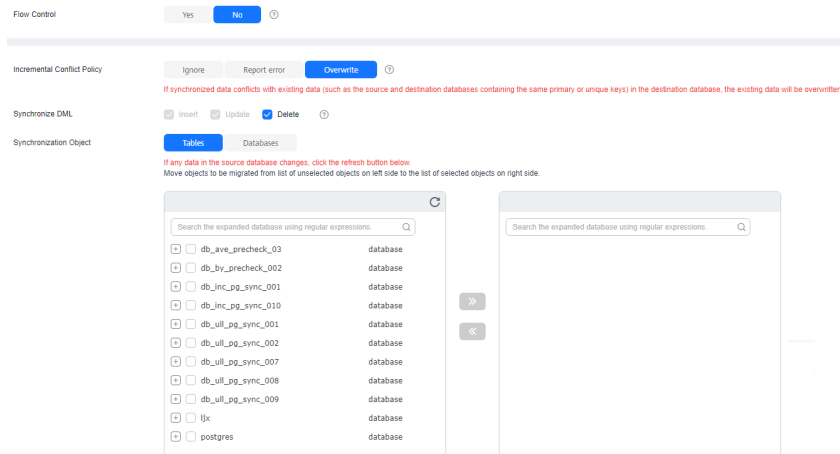
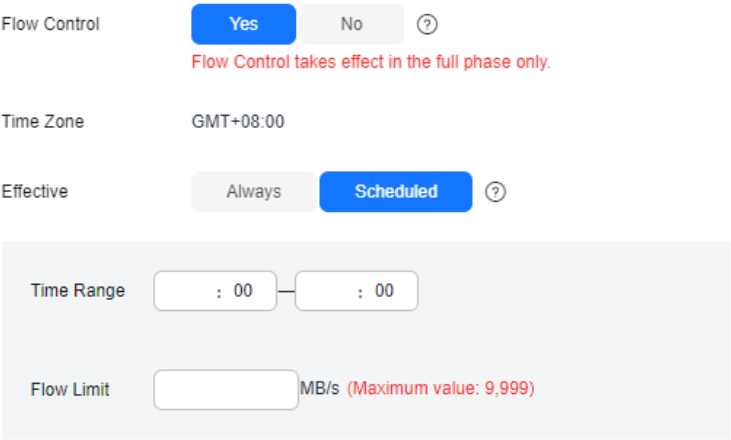



Table 5-245 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-195 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). When you map a schema name or table name to the target object, name the indexes and constraints of the mapped table in the following format: prefix + full name-based hash value + original index/constraint name + _key to prevent index/constraint name conflicts. The prefix of the index is i_, the prefix of the constraint is c_, and the full name is "schema name_table name_index/constraint name". The original index/constraint name may be truncated due to length limitation. You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-196 Task startup settings

Table 5-246 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.22 From GaussDB Distributed to PostgreSQL

Supported Source and Destination Databases

Table 5-247 Supported databases

Source DB	Destination DB
GaussDB Distributed	<ul style="list-style-type: none">• On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)• ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 5-248 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-248 Supported synchronization objects

Type	Constraints
Synchronization scope	<ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● The following field types are not supported: SMALLDATETIME, REETIME, ABSTIME, TID, XID, CID and OID ● Table-level synchronization is supported. <ul style="list-style-type: none"> - The table data can be synchronized. - During full synchronization, only data in the selected tables can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, temporary tables, and tables containing generated columns cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric and rdsRepl) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - Sequence values cannot be synchronized. If there are associated sequences in the table to be synchronized, you can manually synchronize sequence values by referring to Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database after the task is complete. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-249](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 5-249 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see 	<p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB.

Type	Full	Incremental	Full+Incremental
		Connection and Port Description for Incremental Synchronization from GaussDB.	
Destination database user	The user must have the following minimum permissions: <ul style="list-style-type: none"> • Database permission: CONNECT • Schema permission: USAGE • Table permission: INSERT, UPDATE, DELETE, and SELECT • Sequence permission: UPDATE 		

Precautions

The full+incremental synchronization consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full or incremental synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-250 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database has sufficient disk space. - The destination database, schemas, and table object structures must be created in advance. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Constraints
	<ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - During real-time synchronization, you cannot add a coordinator node or data node to the source database. Otherwise, the task fails or data is inconsistent. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - Source database version earlier than 94563 and GaussDB B- or MySQL-compatible mode: If b_format is enabled (for example, set b_format_version = 's1', set b_format_dev_version = 's1', or set b_format_dev_version = 's2'), TIMESTAMP WITH TIME ZONE data cannot be synchronized. To query the source database version, run the following SQL statement: <pre>select working_version_num();</pre>
Full synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Restricted by the GaussDB logical replication function, DDL statements cannot be synchronized. If you synchronize DDL statements, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements.

Type	Constraints
Incremental synchronization	<ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide.
Synchronization comparison	<ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● When GaussDB is not compatible with PostgreSQL, the processing precision of the date data type in the GaussDB database may be different from that in the PostgreSQL database. As a result, data of the two databases is inconsistent.

Type	Constraints
Stopping a task	<p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page.

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-197 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region:

Project:

* Task Name: ⓘ

Description: ⓘ

0/256

Table 5-251 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-198 Synchronization instance details

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

- Synchronization Type: One-way Two-way
- Data Flow: To the cloud Out of the cloud Self-built to self-built
- Source DB Engine: MySQL Oracle DR2 for LUW GaussDB V1R3 **GaussDB Distributed** GaussDB Primary/Standby PingPongSQL Microsoft SQL Server
- Destination DB Engine: MySQL Oracle GaussDB Distributed GaussDB Primary/Standby Kafka **PingPongSQL**
- Network Type: Public network
- VPC: DRS will automatically bind the specified EIP to the DRS instance and order the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service. vpc-2662172-16-0-9123 View VPC
- Synchronization Instance Subnet: subnet-2668172-16-0-9241 View Subnets
- IP Address Type: IPv4 IPv6
- Security Group: default
- Synchronization Mode: Full Incremental Full Incremental
- Source DN Quantity: 2
- Specify EIP: Create an EIP

Table 5-252 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .

Parameter	Description
Destination DB Engine	Select PostgreSQL .
Network Type	<p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Source DN Quantity	The value must be the same as the number of DNs in the distributed source database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-199 AZ



Table 5-253 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-200 Enterprise Project and Tags




Table 5-254 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-201 Source database information

Configure Your Own DNS Server ⓘ

DNS Server

Source Database

CN IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

DN IP Address or Domain Name ⓘ IP Address or Domain Name ⓘ

This button is available only after the replication instance is created successfully.

Table 5-255 Source database settings

Parameter	Description
CN IP Address or Domain Name	Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
DN IP Address or Domain Name	You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different DB instances are different. You are advised to contact O&M personnel.

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-202 Destination database information

Destination Database

Select Connection

VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Name

Database Username

Database Password

SSL Connection

If you want to enable SSL connection, ensure that SSL has been enabled on the destination database, and related parameters have been correctly configured.

This button is available only after the replication instance is created successfully.

Table 5-256 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Name	Indicates whether to specify a database. If this option is enabled, enter the database name.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-203 Synchronization mode

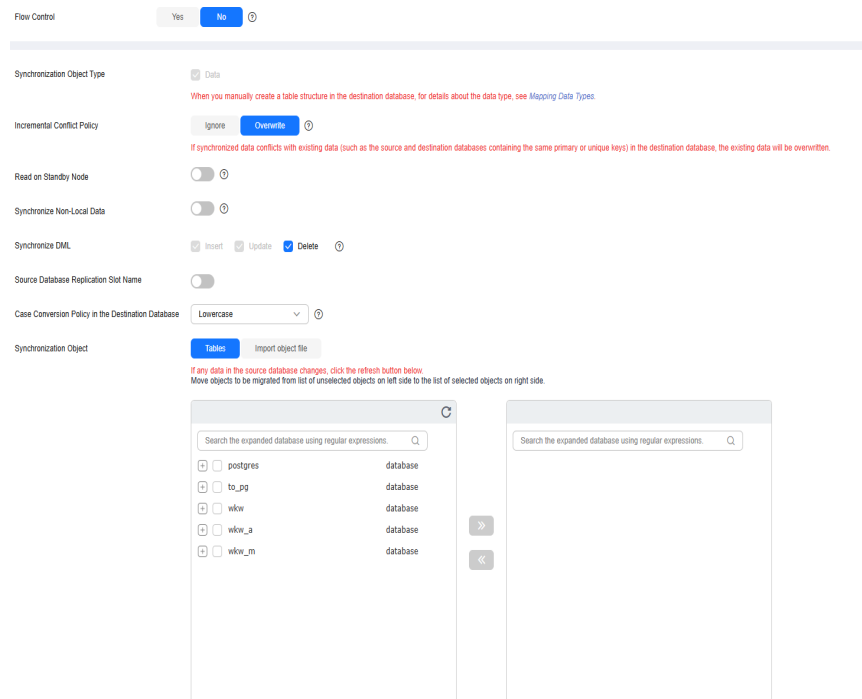
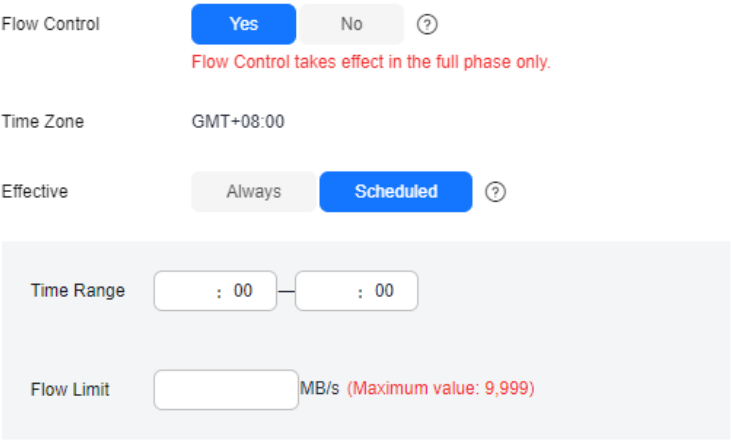



Table 5-257 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 5-204 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	Type of objects for full synchronization. Data is mandatory.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p>
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Source Database Replication Slot Name	<p>You can choose whether to specify the replication slot of the source database. After replication slot is enabled, enter the replication slot name. The name contains 63 characters and cannot start with a digit. Only lowercase letters, digits, and underscores (_) are allowed.</p>
Case Conversion Policy in the Destination Database	<p>Case of schema names, table names, and column names can be converted. If you have specified the mapping name when selecting objects to be synchronized, ignore this parameter.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • You can change object names when you select Import object file. For details, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-205 Task startup settings





Table 5-258 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.23 From DB2 for LUW to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-259 Supported databases

Source DB	Destination DB
DB2 for LUW 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB primary/standby

Precautions

Before creating a synchronization task, read the restrictions listed in the following table.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-260 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • Source database <ul style="list-style-type: none"> - To perform full synchronization, the CONNECT and DATAACCESS permissions are required. - To perform full+incremental synchronization, the DBADM permission is required. - If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the Datacapture attribute of the table to be synchronized is Y. • Destination database <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● User-defined types are not supported. ● The maximum precision supported by the TIMESTAMP type is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency.
Source database	<ul style="list-style-type: none"> ● Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. ● Only the GBK and UTF8 character sets can be synchronized.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• The destination database must be a GaussDB primary/standby instance.• Ensure that a database named in lowercase letters has been created in destination database.• Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.• The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Do not use an imprecise value type as the primary key in the database because the synchronization of UPDATE and DELETE statements in the DRS incremental scenario will be affected. ● Arm VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● Before starting a full+incremental task, ensure that the source database does not contain uncommitted transactions. You are advised to stop writing data to the source database and then start the task. ● After a task is started, the detach operation on a partition table may cause data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot contain uncommitted DDL transactions. ● Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. ● During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. ● During table structure synchronization in the full phase, if there is a schema with the same name as a user in the

Type	Restrictions
	<p>destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. You can add additional objects in the incremental synchronization phase.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-206 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name

Description

0/256

Table 5-261 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

Parameter	Description
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-207 Synchronization instance details

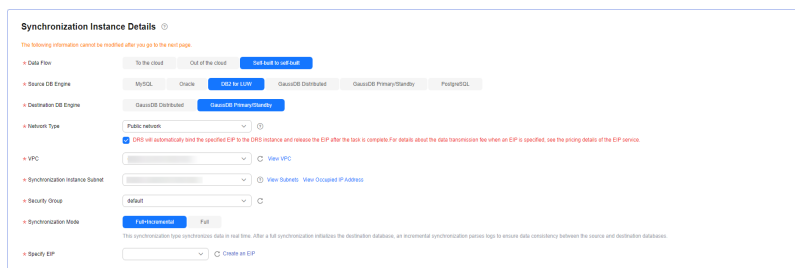


Table 5-262 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select DB2 for LUW.
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-208 AZ

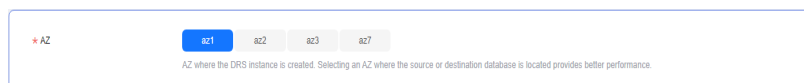


Table 5-263 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-209 Enterprise Project and Tags




Table 5-264 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-210 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized--

This button is available only after the replication instance is created successfully.

Table 5-265 Source database settings


Parameter	Description
IP Address or Domain Name	IP address or domain name of the source DB2 for LUW database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source DB2 for LUW database.
Database Password	The password for the source database username.
Database Name	The name of the database to which the synchronization objects belong in the source DB2 for LUW.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures.
Synchronize Driver	You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection.

 **NOTE**

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-211 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-266 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-212 Synchronization mode

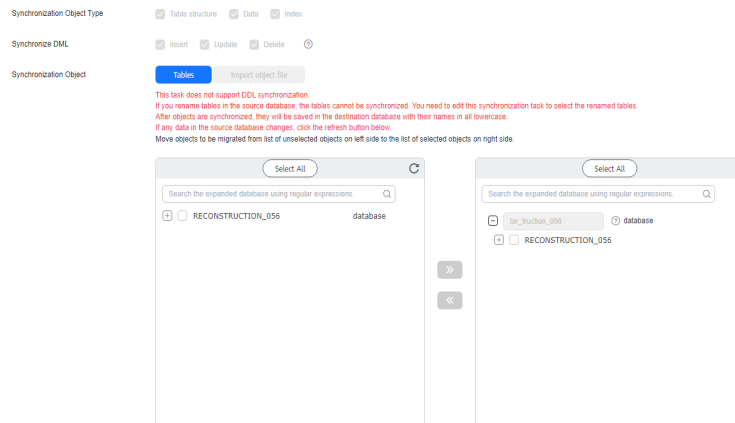



Table 5-267 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can synchronize tables or import object files based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-213 Task startup settings

Table 5-268 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.24 From DB2 for LUW to GaussDB Distributed

Supported Source and Destination Databases

Table 5-269 Supported databases

Source DB	Destination DB
DB2 for LUW 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB distributed

Precautions

Before creating a synchronization task, read the restrictions listed in the following table.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-270 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database <ul style="list-style-type: none"> - To perform full synchronization, the CONNECT and DATAACCESS permissions are required. - To perform full+incremental synchronization, the DBADM permission is required. - If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the Datacapture attribute of the table to be synchronized is Y. ● Destination database <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● User-defined types are not supported. ● The maximum precision supported by the TIMESTAMP type is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency.
Source database	<ul style="list-style-type: none"> ● Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. ● Only the GBK and UTF8 character sets can be synchronized.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• The destination database must be a GaussDB distributed instance.• Ensure that a database named in lowercase letters has been created in destination database.• Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.• The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Do not use an imprecise value type as the primary key in the database because the synchronization of UPDATE and DELETE statements in the DRS incremental scenario will be affected. ● Arm VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● Before starting a full+incremental task, ensure that the source database does not contain uncommitted transactions. You are advised to stop writing data to the source database and then start the task. ● After a task is started, the detach operation on a partition table may cause data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot contain uncommitted DDL transactions. ● Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. ● During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. ● During table structure synchronization in the full phase, if there is a schema with the same name as a user in the

Type	Restrictions
	<p>destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created.</p> <ul style="list-style-type: none"> • Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • You can add additional objects in the incremental synchronization phase. • During an incremental synchronization, if you update the primary key column or the first column of a table that does not have a primary key in the source database DB2 for LUW, the distribution column will be updated in the GaussDB database, which may cause data inconsistency. Thus, do not update the preceding columns.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-214 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 5-271 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-215 Synchronization instance details

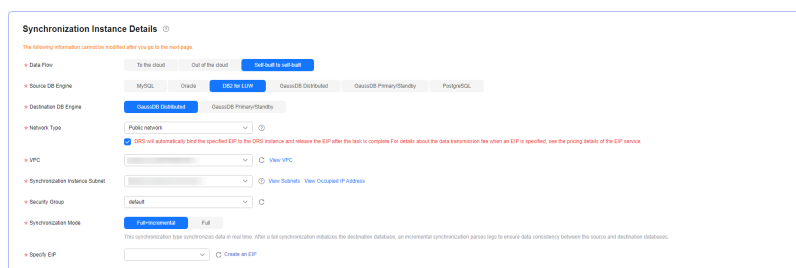


Table 5-272 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select DB2 for LUW.
Destination DB Engine	Select GaussDB Distributed .
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-216 AZ




Table 5-273 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-217 Enterprise Project and Tags




Table 5-274 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-218 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver

This button is available only after the replication instance is created successfully.

Table 5-275 Source database settings


Parameter	Description
IP Address or Domain Name	IP address or domain name of the source DB2 for LUW database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source DB2 for LUW database.
Database Password	The password for the source database username.
Database Name	The name of the database to which the synchronization objects belong in the source DB2 for LUW.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> • The maximum size of a single certificate file that can be uploaded is 500 KB. • If the SSL certificate is not used, your data may be at risk. • Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth • DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures.
Synchronize Driver	You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection.

 **NOTE**

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-219 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username


Database Password 

Table 5-276 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). Example: 192.168.0.1:8000,192.168.0.2:8000
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-220 Synchronization mode

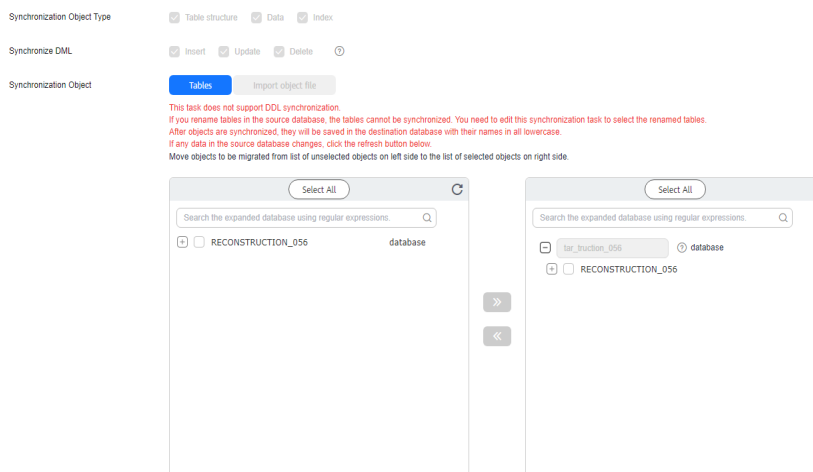



Table 5-277 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Changing Object Names (Mapping Object Names). • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-221 Task startup settings

Table 5-278 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

5.25 From Microsoft SQL Server to Kafka

Supported Source and Destination Databases

Table 5-279 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) 	Kafka 0.11 or later

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 5-280 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-280 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY Scope of incremental synchronization <ul style="list-style-type: none"> DML statements, including INSERT, UPDATE, and DELETE, are supported. DDL statements are not supported.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-281](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-281 Database account permission

Type	Incremental
Source database user	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized

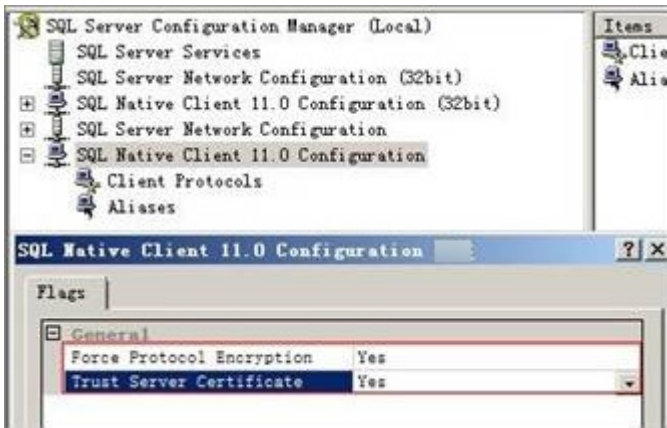
Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-282 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - The SQL Server Agent proxy service must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 5-222. <p style="text-align: center;">Figure 5-222 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database parameter requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later.

Type	Restrictions
	<ul style="list-style-type: none"> – Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.
Incremental synchronization	<ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • To ensure normal synchronization and data consistency, you are advised to create a DRS task again during off-peak hours to meet the preceding requirements. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • DDL operations performed on the source database will not be synchronized to the destination database. • The IMAGE, TEXT, and NTEXT big data types cannot be deleted. • You can add additional synchronization objects.

Procedure

This section uses self-managed Microsoft SQL Server to Kafka synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-223 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click **Create Now** on this page.
The system will create virtual resources immediately after you click **Create Now**. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 5-283 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 5-224 Synchronization instance details

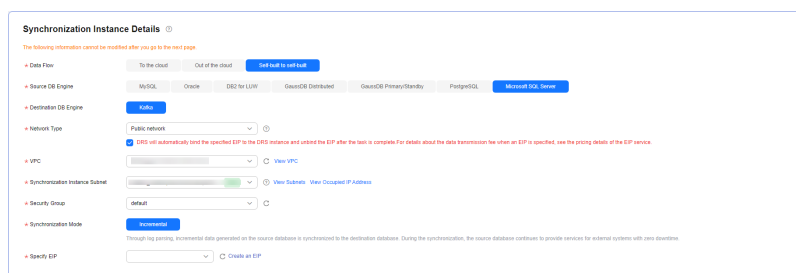


Table 5-284 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select Kafka .

Parameter	Description
Network Type	<p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 5-225 AZ




Table 5-285 Task AZ

Parameter	Description
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 5-226 Enterprise Project and Tags




Table 5-286 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-227 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 5-287 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source Microsoft SQL Server database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for logging in to the source Microsoft SQL Server database.
Database Password	The password for the database username.

Figure 5-228 Destination database information

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Table 5-288 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 5-229 Synchronization Object

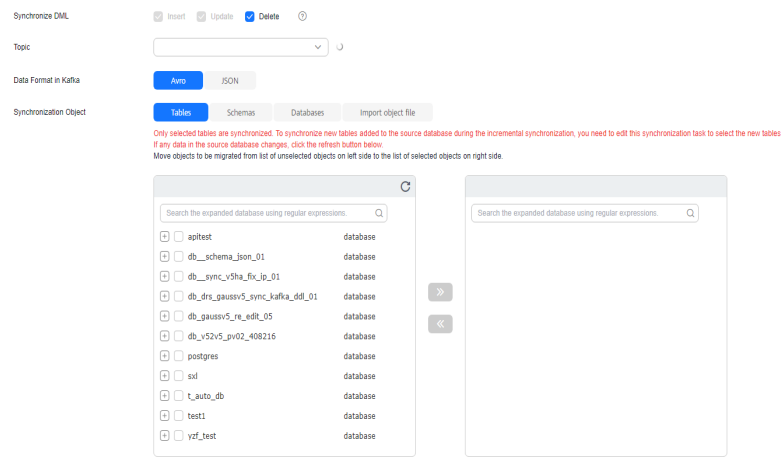



Table 5-289 Synchronization mode and object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected. If you do not select Delete , DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.
Topic	Select the topic to be synchronized to the destination database.
Data Format in Kafka	Select the format of data delivered to Kafka. <ul style="list-style-type: none"> ● Avro refers to binary encoded format. ● Json refers to data interchange format. For details, see Kafka Message Format .

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. For details about the import procedure and description, Importing Synchronization Objects. • When you select Import object file, you can use the mapping function in Changing Object Names (Mapping Object Names) only when the topic synchronization policy is set to A specific topic. Otherwise, topics are generated based on the name format. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-230 Task startup settings





Table 5-290 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .
Delay Threshold (s)	During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications.

Parameter	Description
Data Exception Notification	This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

6 Two-Way Synchronization

6.1 Two-Way Synchronization from MySQL to MySQL

Supported Source and Destination Databases

Table 6-1 Supported databases

Instance Role in Current Cloud	Source DB	Destination DB
Active 1	RDS for MySQL (5.5, 5.6, 5.7, and 8.0)	<ul style="list-style-type: none">On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)RDS for MySQL (5.5, 5.6, 5.7, and 8.0)
Active 2	<ul style="list-style-type: none">On-premises MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)ECS-hosted MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)Other cloud MySQL databases (MySQL 5.5, 5.6, 5.7, and 8.0)RDS for MySQL (5.5, 5.6, 5.7, and 8.0)	RDS for MySQL (5.5, 5.6, 5.7, and 8.0)

 NOTE

- The version of the destination database must be the same as that of the source database.
- Only whitelisted users can use this function. To use this function, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket**.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 6-2](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 6-2 Database account permissions

Type	Forward Task	Backward Task
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT If the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup or task editing.	The user must have the following minimum permissions: REPLICATION SLAVE and REPLICATION CLIENT
Destination database user	The user must have the following minimum permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES The root account of the RDS for MySQL DB instance has the preceding permissions by default. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 6-3 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 6-3 Supported synchronization objects

Type	Notes
Synchronization objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre>

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see **Pre-check Items**. In addition to the pre-check items, you need to pay attention to the items listed in **Table 6-4**.

Table 6-4 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • The source and destination databases cannot contain tables that have the same names but do not have primary keys. • If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● You are not advised to modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. ● Do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● Do not clear binlogs on the source and destination databases. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. ● During data synchronization, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the synchronization task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a synchronization task again if the source MySQL database is upgraded across major versions. <p>Full synchronization</p> <ul style="list-style-type: none"> ● When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● Modifying MyISAM tables may cause data inconsistency. <p>Incremental synchronization</p> <ul style="list-style-type: none"> ● If the session variable character_set_client is set to binary, some data may include garbled characters. ● You can add additional objects in the incremental synchronization phase of a forward task. ● Resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down. ● Resetting a task is not supported for a two-way synchronization task. ● During incremental synchronization of a forward task, some DDL statements are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and

Type	Restrictions
	<p>ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required.</p> <ul style="list-style-type: none"> - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • DDL synchronization is not supported during incremental synchronization of a backward task. • Do not perform the point-in-time recovery (PITR) operation on the source database. • The destination database cannot be restored to a point in time when a full synchronization was being performed. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization logs. • Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p>Troubleshooting</p> <ul style="list-style-type: none"> • If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB of memory or higher instance specifications, the synchronization task cannot be created. ● During table name mapping, tables on which views, stored procedures, and functions depend cannot be synchronized, and foreign key constraints of tables cannot be synchronized. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the source and destination databases are the same RDS DB instance, real-time synchronization without database mapping is not supported. ● If the database table name contains characters other than letters, digits, and underscores (<code>_</code>), or the mapped database table name contains hyphens (<code>-</code>) and number signs (<code>#</code>), the name length cannot exceed 42 characters. ● Braces (<code>{}</code>) in a database or table name in the source database must be used in pairs. Otherwise, the table structure synchronization may fail. After data mapping, if braces (<code>{}</code>) in a database or name mapped to the destination database are not in pairs, the table structure may fail to be synchronized. ● The partitioned table does not support column mapping. ● If the source database contains non-standard floating-point data and the data can be written in loose mode but cannot be written in strict mode, there may be data inconsistency during synchronization. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● When creating a task, ensure that the synchronization object dimensions (table-level, database-level, and object file import) selected for the forward task and backward task are the same. ● After a synchronization task is created, the source and destination databases cannot be set to read-only. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted

Type	Restrictions
	<p>data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. • For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed. Otherwise, all synchronization tasks fail. • The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. • Do not perform service operations on the same data record in the same table in the source database of the forward task and the source database of the backward task at the same time. Otherwise, data may be inconsistent.

Prerequisites

- [You have logged in to the DRS console.](#)
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases.](#)
- If a subaccount is used to create a DRS task, ensure that an agency has been added. For details about how to create an agency, see [Agency Management.](#)

Procedure

This section describes how to create a two-way synchronization task from MySQL to RDS for MySQL.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 6-1 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu] Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

Task Name: [Text input: DRS-5678]

Description: [Text area: 0/256]

Table 6-5 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance details

Figure 6-2 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

- Synchronization Type: One-way Two-way
- Instance Role in Current Cloud: Active 1 Active 2
Active 2 indicates that the DB instance serves as the destination database of the forward task and is waiting for data synchronization from the source database. It also serves as the source database of the backward task.
- Source DB Engine: MySQL
- Destination DB Engine: MySQL
- Network Type: Public network
 DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.
- Destination DB Instance: Select an instance View DB instance View Unavailable DB Instance
During the full synchronization of a DRS task, a lot of logs are generated. These logs may be incorrectly stored locally, which may cause the storage space to be used up. You are advised to enable storage autoresizing for the RDS DB instance. During the DRS task, set an appropriate local retention period for RDS logs. You can also clear logs exceeding the specified retention period with just a few clicks.
- Synchronization Instance Subnet: Select the subnet View Subnets
- Security Group: default
- Synchronization Mode: Full-synchronous
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.
- Enable Binlog Cleanup:
- Specify EIP: Choose an EIP

Table 6-6 Synchronization instance settings

Parameter	Description
Synchronization Type	Select Two-way . A forward task and a backward task are created. The synchronization mode of the forward task is full+incremental synchronization, and the synchronization mode of the backward task is incremental synchronization.
Instance Role in Current Cloud	Specifies the role of the current cloud DB instance in the two-way synchronization. The value can be Active 1 and Active 2 . For details, see How Do I Select Active Database 1 and 2 for Dual-Active DR? <ul style="list-style-type: none"> - Active 1: Indicates that the DB instance serves as the source database of the forward task and contains initial data. - Active 2: Indicates that the DB instance serves as the destination database of the forward task. Active 2 is used as an example.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination database bound with an EIP. - VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select an RDS DB instance you have created.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Security Group	Select a security group. You can use security group rules to allow or deny access to the instance.
Synchronization Mode	The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental , data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.

- Task type

Figure 6-3 Task type




Table 6-7 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Project and Tags

Figure 6-4 Enterprise Project and Tags




Table 6-8 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, the **Configure Source and Destination Databases** page of the forward task is displayed by default. On the **Configure Source and Destination Databases** page of the forward task, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

The information to be entered for the source and destination databases is different for the role Active 1 and Active 2 of the DB instance on the current cloud.

- Source database information of the forward task

Figure 6-5 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 6-9 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Username	The username for accessing the source database.
Database Password	The password for the database username. You can change the password if necessary. If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.
SSL Connection	If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.


- Destination database information of the forward task

Figure 6-6 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

Table 6-10 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance you selected when creating the synchronization task. This parameter cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username. You can change the password if necessary. If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.
SSL Connection	If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 6-7 Synchronization mode

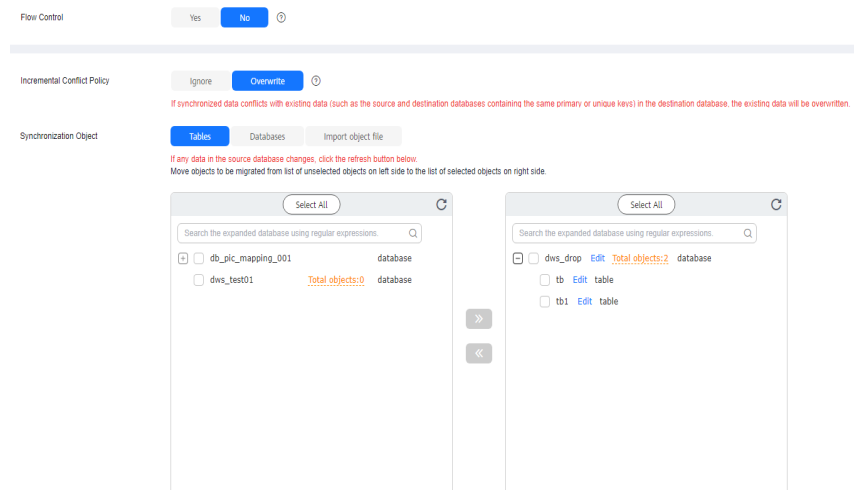
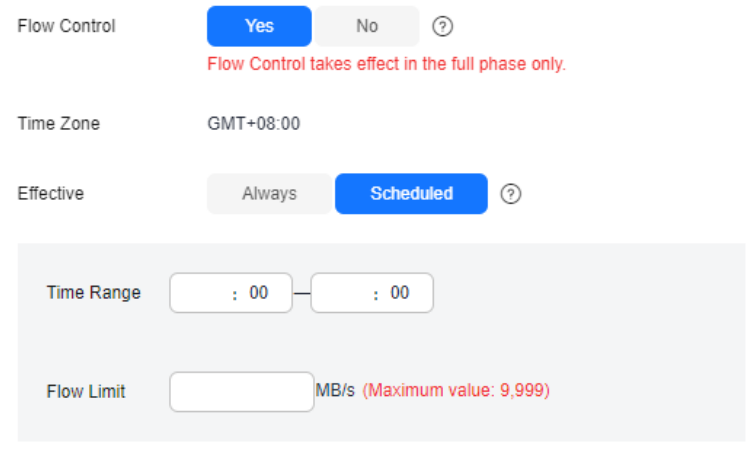



Table 6-11 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 6-8 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> ● If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – In a two-way synchronization scenario, many-to-one database and table name mapping is not supported. – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – If you want to change the saved database mapping name during table-level synchronization, you need to expand the database. ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 6-9 Task startup settings

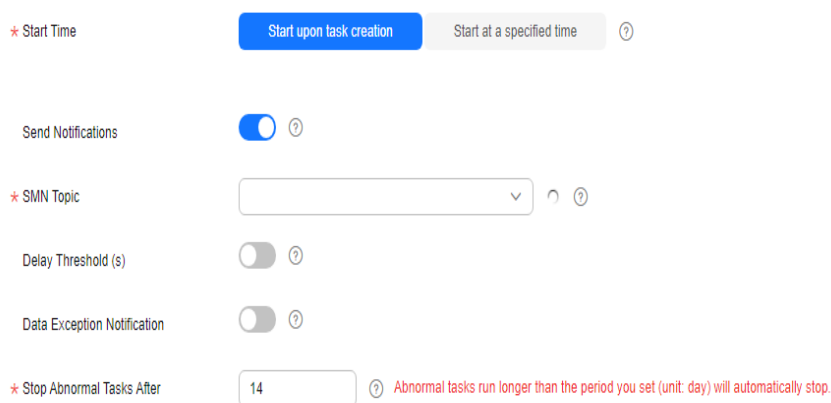


Table 6-12 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 7 After the forward task is submitted, wait until the forward task starts and enters the incremental synchronization state. After the latency is less than 60 seconds, click **Edit** of the backward task on the task list page to test connections of the backward task. After the connection tests are successful, click **Next**.

The connection information of the backward task cannot be edited. You can only test the connections.

Step 8 On the **Set Synchronization Task** page of the backward task, select the conflict policy and synchronization objects, and then click **Next**.

Figure 6-10 Synchronization mode

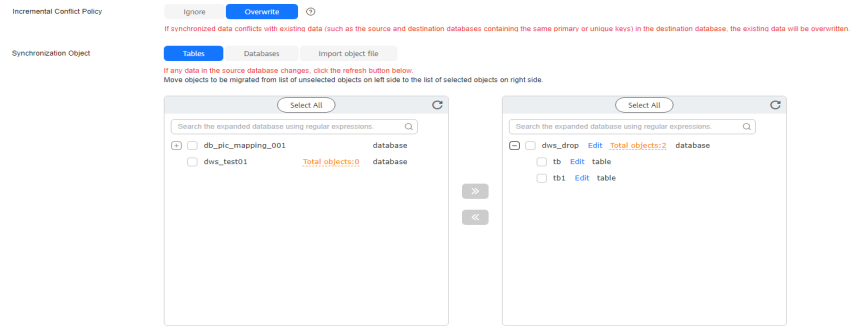



Table 6-13 Synchronization mode and object

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. Overwrite Conflicting data will be overwritten.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • When selecting the synchronization object, ensure that the object dimensions (table-level, object file import, and database-level) selected for the forward task and backward task are the same. The synchronization object of the forward task must be consistent with that of the backward task. For example: If the synchronization object for the forward task is database B (table T1), the synchronization object for the backward task must be database B (table T1). • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> – In a two-way synchronization scenario, many-to-one database and table name mapping is not supported. – In a backward task, the object mapping must be consistent with that in a forward task. For example: If the synchronization object for the forward task is database B (mapped to database B1) [table T1 (mapped to table T2)], the synchronization object for the backward task must be database B1 (mapped to database B) [table T2 (mapped to table T1)]. – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – If you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 9 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.


- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 10 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the backward task.

Step 11 After the backward task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- You can switch over the forward task and backward task after the status of the backward task changes to **Incremental**. For details, see [Exchanging the Direction for a Two-Way Synchronization Task](#).

----End

7 Task Management

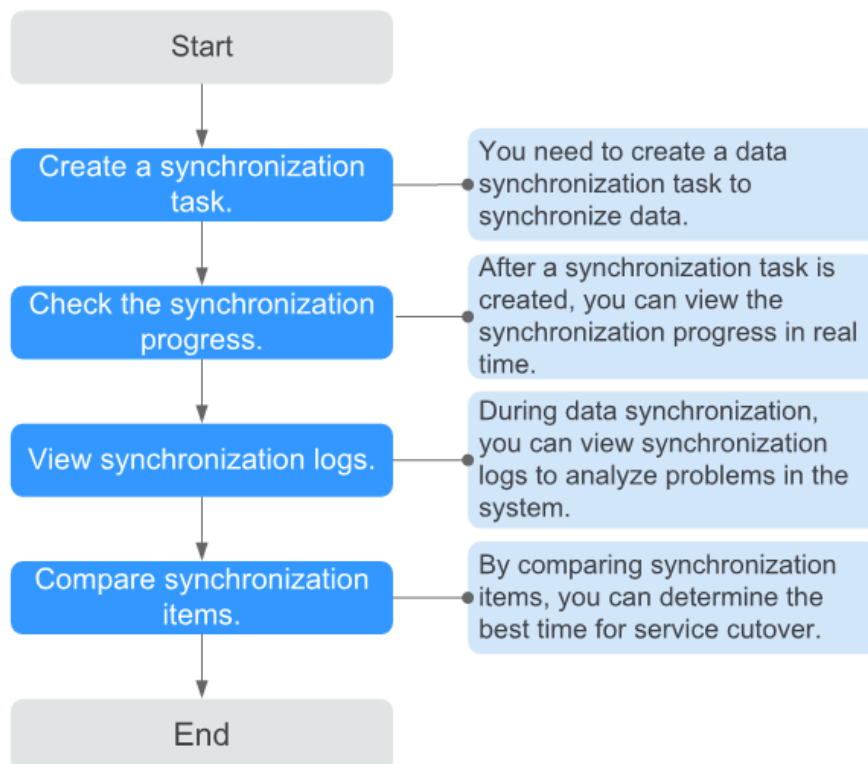
7.1 Creating a Synchronization Task

Process

A complete real-time synchronization consists of creating a synchronization task, tracking task progress, analyzing synchronization logs, and comparing data consistency. By comparing multiple items and data, you can synchronize data between different service systems in real time.

A complete real-time synchronization involves the following procedures.

Figure 7-1 Flowchart



- **Step 1: Create a synchronization task.** Select the source and destination databases as required and create a synchronization task.
- **Step 2: Check the synchronization progress.** During synchronization, you can view the synchronization progress.
- **Step 3: View synchronization logs.** Synchronization logs contain alarms, errors, and prompt information. You can analyze system problems based on such information.
- **Step 4: Compare synchronization items.** You can compare objects and data to be synchronized to ensure data consistency.

This section describes how to synchronize data from a MySQL database to an RDS for MySQL database. To configure other storage engines, you can refer to the following procedures.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 7-2 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 7-1 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region.
Project	The project corresponds to the current region and can be changed.
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\

- Synchronization instance information

Figure 7-3 Synchronization instance information

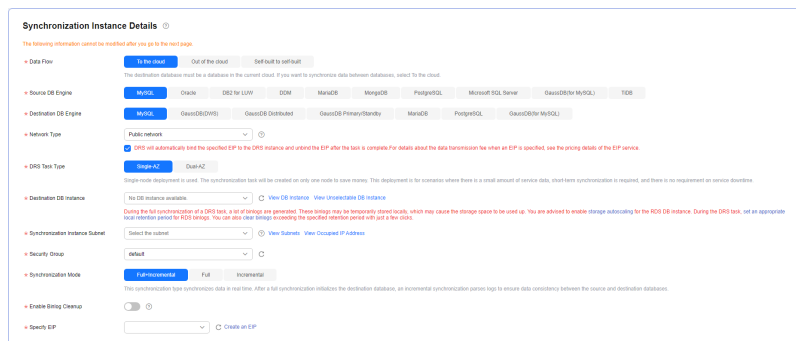


Table 7-2 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is a database in the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	<p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region and VPC. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for data synchronization from on-premises databases to cloud databases, between databases of different accounts in the same region on the cloud, or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.

Parameter	Description
DRS Task Type	<p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p>
Destination DB Instance	<p>The RDS DB instance you created.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The destination DB instance cannot be a read replica. - The source and destination DB instances can be the same DB instance.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Security Group	<p>Select a security group. You can use security group rules to allow or deny access to the instance.</p>

Parameter	Description
Synchronization Mode	<p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full synchronization and disabled during the incremental synchronization.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>If DRS Task Type is set to Dual-AZ, you need to specify the primary and standby IP addresses.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- Task Type

Figure 7-4 Task type





Table 7-3 Task type information

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>
AZ	<p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 7-5 AZ</p> 

- Enterprise Project and Tags

Figure 7-6 Enterprise Project and Tags




Table 7-4 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p>

Parameter	Description
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

In different data flow scenarios, the source and destination database settings are different. Specify the required parameters based on the GUI.

- (Optional) Configuring your own DNS server

Figure 7-7 DNS Server




Table 7-5 DNS server information

Parameter	Description
DNS Server	Enable this option if you need to use the IP address of your own DNS server as the source or destination database IP address.
DNS Server IP Address	Add the IP address of your own DNS server to DNS Server IP Address . Then, you can also enter this IP address in IP Address or Domain Name in the Source Database or Destination Database area for data migration.

 **NOTE**

This function is available when you need to use the IP address of your own DNS server as the source or destination database IP address.

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- Source database information

Figure 7-8 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

● Test successful

Table 7-6 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.

Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

- Destination database information

Figure 7-9 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

SSL Connection

Table 7-7 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>

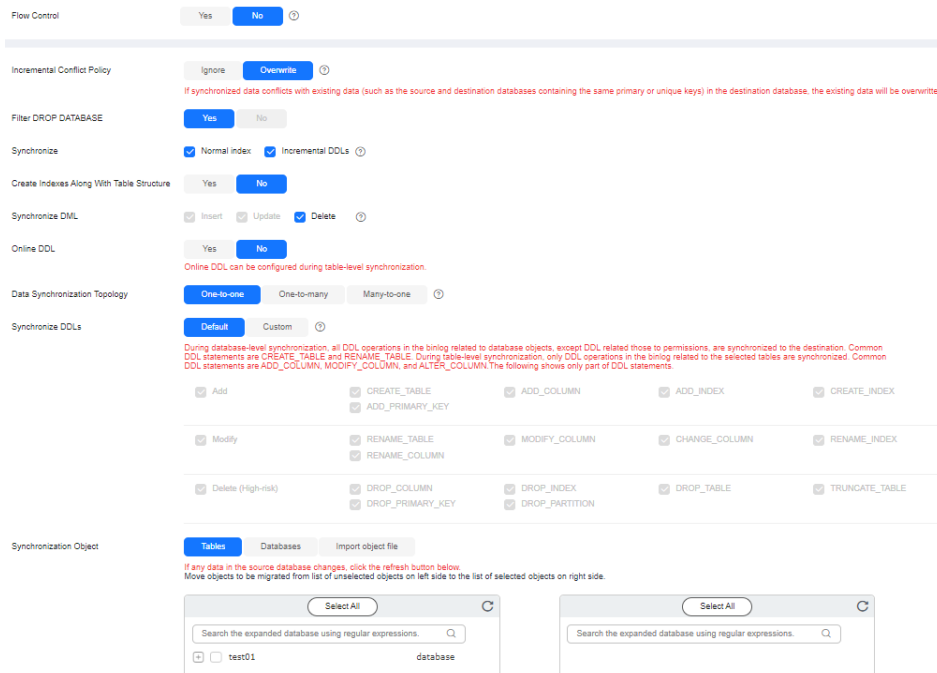
Parameter	Description
SSL Connection	<p>If SSL connection is required, enable SSL on the destination database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 7-10 Synchronization mode



Flow Control: Yes No

Incremental Conflict Policy: Ignore Overwrite
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Filter DROP DATABASE: Yes No

Synchronize: Normal Index Incremental DDLs

Create Indexes Along With Table Structure: Yes No

Synchronize DML: Insert Update Delete

Online DDL: Yes No
Online DDL can be configured during table-level synchronization.

Data Synchronization Topology: One-to-one One-to-many Many-to-one

Synchronize DDLs: Default Custom
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. The following shows only part of DDL statements.

<input checked="" type="checkbox"/> Add	<input checked="" type="checkbox"/> CREATE_TABLE <input checked="" type="checkbox"/> ADD_PRIMARY_KEY	<input checked="" type="checkbox"/> ADD_COLUMN	<input checked="" type="checkbox"/> ADD_INDEX	<input checked="" type="checkbox"/> CREATE_INDEX
<input checked="" type="checkbox"/> Modify	<input checked="" type="checkbox"/> RENAME_TABLE <input checked="" type="checkbox"/> RENAME_COLUMN	<input checked="" type="checkbox"/> MODIFY_COLUMN	<input checked="" type="checkbox"/> CHANGE_COLUMN	<input checked="" type="checkbox"/> RENAME_INDEX
<input checked="" type="checkbox"/> Delete (High-risk)	<input checked="" type="checkbox"/> DROP_COLUMN <input checked="" type="checkbox"/> DROP_PRIMARY_KEY	<input checked="" type="checkbox"/> DROP_INDEX <input checked="" type="checkbox"/> DROP_PARTITION	<input checked="" type="checkbox"/> DROP_TABLE	<input checked="" type="checkbox"/> TRUNCATE_TABLE

Synchronization Object: Tables Databases Import object file
If any data in the source database changes, click the refresh button below. Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

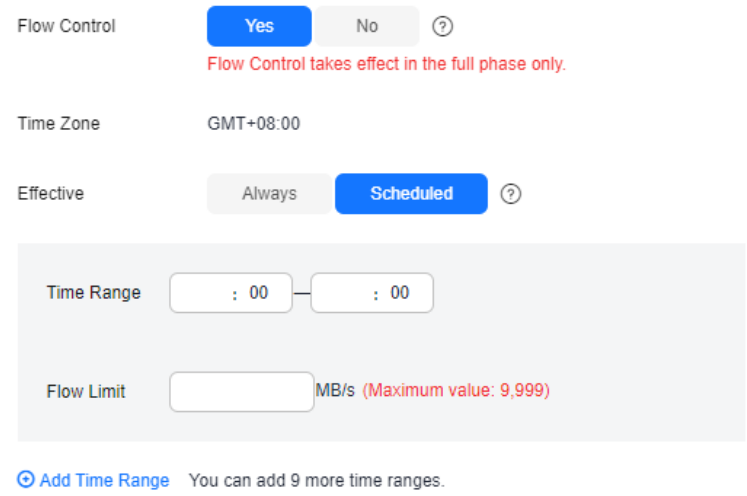
Search the expanded database using regular expressions.

test01 database

Select All

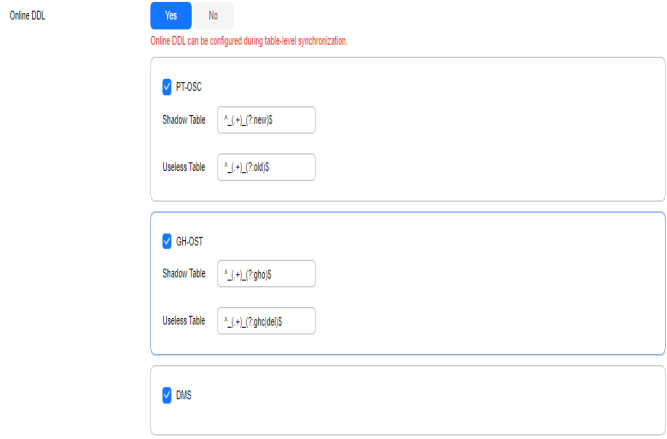
Search the expanded database using regular expressions.

Table 7-8 Synchronization mode and object


Parameter	Description
Flow Control	<p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. During the full synchronization, the synchronization speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 7-11 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize normal indexes and DDLs based on service requirements.</p>
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full synchronization phase.</p> <ul style="list-style-type: none"> ● Yes: Indexes are migrated when the table structure is migrated in the full synchronization phase. ● No: Indexes are migrated separately after data synchronization.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> <p>If you do not select Delete, DELETE statements in the incremental data of the source database will not be synchronized, which may cause a data inconsistency. As a result, there may be a data conflict or the task may fail.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> <p>For details, see How Do I Specify the Start Point for DRS Incremental Synchronization?</p>

Parameter	Description
Data Synchronization Topology	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p>
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: none"> - During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. - During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables.

Parameter	Description
Online DDL	<p>If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.</p> <ul style="list-style-type: none"> Yes: Table-level synchronization supports Online DDL synchronization. You can select three Online DDL tools including PT-OSC, GH-OST, and DMS. The regular expressions can be configured for shadow tables and useless tables in PT-OSC and GH-OST. (Default values are provided. You are advised not to change the regular expressions unless otherwise specified.) <p>Figure 7-12 Online DDL</p>  <p>The screenshot shows the 'Online DDL' configuration page. At the top, there are 'Yes' and 'No' buttons, with 'Yes' selected. Below this, a red note states 'Online DDL can be configured during table-level synchronization'. There are three main configuration blocks:</p> <ul style="list-style-type: none"> PT-OSC: A checked checkbox. Below it are two input fields: 'Shadow Table' with the value '^_+_?(?mev)\$' and 'Useless Table' with the value '^_+_?(?old)\$'. GH-OST: A checked checkbox. Below it are two input fields: 'Shadow Table' with the value '^_+_?(?gho)\$' and 'Useless Table' with the value '^_+_?(?ghode)\$'. DMS: A checked checkbox with no further input fields.

- No:** Table-level synchronization does not support Online DDL synchronization.

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Changing Object Names (Mapping Object Names). <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 7-13 Processing data



Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
gblest_order_payment_1	gblest_order_payment_1	c1	The serverName@database@table column is used	varchar(191)	Add Delete
gblest_order_payment_10	gblest_order_payment_10	c2	Default 1214	int	Add Delete
gblest_order_payment_11	gblest_order_payment_11	--	--	--	Add

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 7-14 Task startup settings

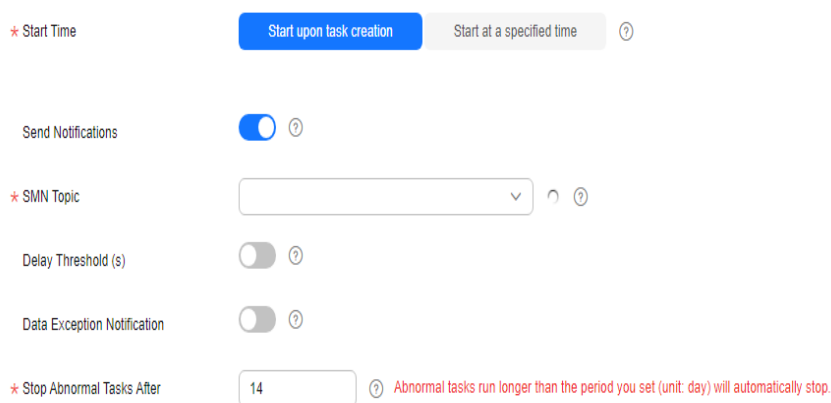



Table 7-9 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide .

Parameter	Description
Delay Threshold (s)	<p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications.
Data Exception Notification	<p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p>
Stop Abnormal Tasks After	<p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

Helpful Links

- [Supported Databases](#)

- [Preparations](#)
- [Synchronization Overview](#)
- [Data Synchronization Topologies](#)

7.2 Querying the Synchronization Progress

This section describes how to check the synchronization progress of each synchronization object.

- During the full synchronization, DRS displays the progress overview. You can view the synchronization progress of structures, data, and indexes. The synchronization progress is displayed based on the number of synchronized objects. When the progress reaches 100%, the synchronization is complete. The synchronization of data and indexes is relatively slow.
- During the full synchronization, you can view the synchronization progress of a specific object and the synchronization details. If the number of objects is the same as that of synchronized objects, the synchronization is complete. You can view the synchronization progress of each object in detail. During incremental synchronization, the progress details are not displayed. You can view the consistency status on the **Synchronization Comparison** tab.
- During an incremental synchronization, DRS displays the incremental synchronization delay. You can determine the synchronization status between the source and destination databases based on the delay. If the delay is 0, the source and destination databases are instantaneously consistent, and no new transaction needs to be synchronized.

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click **Synchronization Progress** to view table synchronization progress.
 - Check the synchronization progress in percentage.
 - When a full synchronization is complete, the progress reaches 100%.
 - After the full synchronization is complete, the incremental synchronization starts. You can view the incremental synchronization delay on the **Synchronization Progress** tab.
 - You can also view the incremental synchronization delay on the **Data Synchronization Management** page. When the incremental synchronization delay exceeds the preset or default threshold, the value of the incremental synchronization delay is displayed in red in the task list.
 - When the delay is 0s, data in the source database is synchronized to the destination database in real time.

 **NOTE**

"Delay" refers to the delay from when the transaction was submitted to the source database to when it is synchronized to the destination database and executed.

Transactions are synchronized as follows:

1. Data is extracted from the source database.
2. The data is transmitted over the network.
3. DRS parses the source logs.
4. The transaction is executed on the destination database.

If the delay is 0, the source database is consistent with the destination database, and no new transactions need to be synchronized.

 **CAUTION**

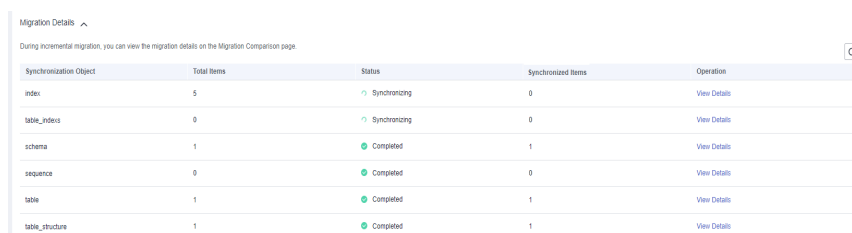
Transactions that have been there for a long time without being submitted and frequent DDL operations may result in excessive synchronization delay.


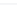
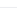
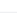
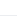

- View the synchronization progress details. In the **Migration Details** area, locate the target synchronization object and click **View Details** in the **Operation** column to view the synchronization progress.

 **NOTE**

The synchronization progress details only display the data synchronized in the full synchronization and the skipped structures and cannot be updated. The progress details are not displayed after incremental synchronization starts. You can compare synchronization items on the **Synchronization Comparison** page.

Figure 7-15 Synchronization progress details



Synchronization Object	Total Items	Status	Synchronized Items	Operation
index	5	 Synchronizing	0	View Details
table_index	0	 Synchronizing	0	View Details
schema	1	 Completed	1	View Details
sequence	0	 Completed	0	View Details
table	1	 Completed	1	View Details
table_structure	1	 Completed	1	View Details

- Skip synchronization objects. In the **Migration Details** area, locate the synchronization object and click **View Details** in the **Operation** column. On the displayed page, click **Skip and Restart** in the **Operation** column to skip the object.

NOTE


- For a synchronization task from MySQL, MariaDB, GaussDB(for MySQL), or TiDB, synchronization objects can be skipped during the full synchronization.
- Migration objects that can be skipped include databases, functions, procedures, table_indexes, table_structure, and views.
- After a synchronization object is skipped, the task will be resumed. If the task is paused or abnormal, the task will be restarted after you click **Skip and Restart**.
- If a task is in the full synchronization state, an object may be synchronized before being skipped. In this case, skipping the object is not applied.
- After an object is skipped, if operations associated with the object are still performed during subsequent synchronization, the task fails. For example, after table_structure is skipped, the task fails because the object is not found when indexes and data of the table are synchronized.
- After an object is skipped, the result for a comparison is inconsistent.
- Skipping table_structure of a table without a primary key will cause data inconsistency. In addition, the task may fail because the table is not found. Perform the operation after confirmation.
- If a DRS task fails because the table is not found after you skip a table without a primary key, contact O&M engineers for the destination database to create the corresponding table in the destination database based on the table structure of the source database and error log information, and click **Resume** in the **Operation** column of the task to submit the task again.

Step 3 In the MySQL synchronization scenario, you can view the information about the source position and consumed position on the **Synchronization Progress** tab.

NOTE

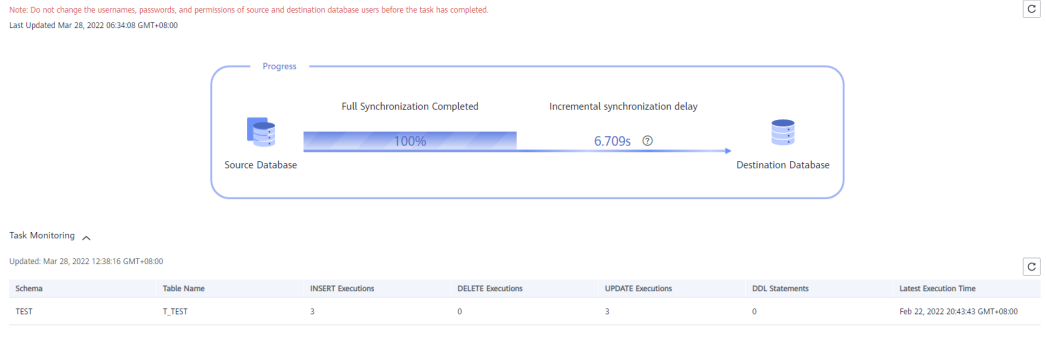
The displayed position information is updated every 10 seconds.

Figure 7-16 Synchronization position information



Step 4 In synchronization with Oracle serving as the source, MySQL to Kafka synchronization, and GaussDB to Kafka synchronization, on the **Synchronization Progress** tab, view the number of DML operations (Insert, Delete, and Update) performed on the source database. In the upper right corner of the **Task Monitoring** list, refresh the list and view the latest monitoring data.

Figure 7-17 Task monitoring

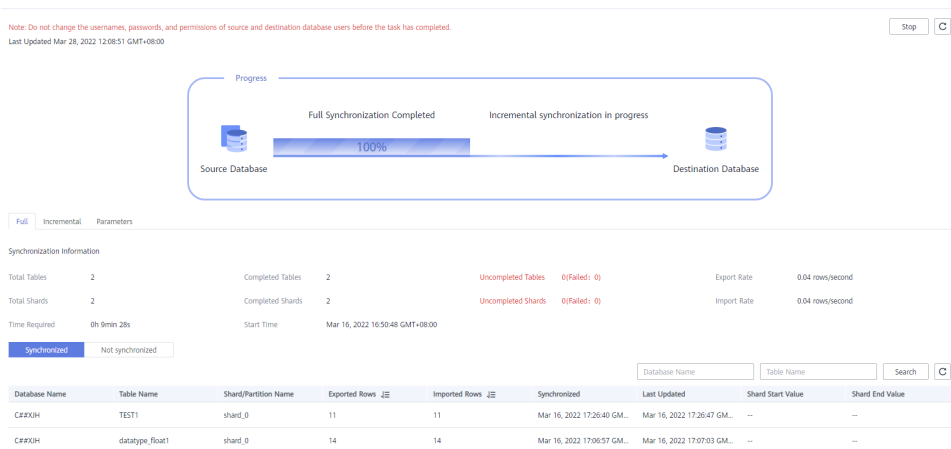


NOTE

1. After you perform the DML operation on the source database and run the commit command to make the operation take effect, the operation data can be displayed in the **Task Monitoring** list.
2. DRS collects statistics on the number of operations based on redo logs.
3. For tables with large object (LOB) data type columns, Oracle records more information in redo logs than the information generated by the actual operations that is performed. The INSERT and UPDATE operations are first performed on non-LOB columns and then on LOB columns. DRS collects statistics only from the redo logs, so the number of recorded operations may be inconsistent with the actual number of changed rows according to operation audit.
4. The Oracle MERGE statement can be converted into INSERT, UPDATE, and DELETE operations in the redo log. Thus, the number of rows changed by the MERGE statement are increased because the INSERT, UPDATE, and DELETE operations are separately recorded in the redo log.

Step 5 In the Oracle to GaussDB primary/standby or distributed, GaussDB primary/standby to MySQL, GaussDB distributed to MySQL, Oracle to MySQL (incremental), and Oracle to GaussDB(DWS) (incremental) synchronization scenarios, search for the detailed synchronization object information on the **Synchronization Progress** tab. The keywords for searching synchronization objects are case sensitive.

Figure 7-18 Synchronization details



----End

7.3 Viewing Synchronization Logs

Synchronization logs refer to the warning-, error-, and info-level logs generated during the synchronization process. This section describes how to view synchronization logs to locate and analyze database problems.

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click **Synchronization Logs** to view the logs generated during the synchronization.

You can view time, levels, and descriptions of the logs.

Figure 7-19 Synchronization logs

Time	Level	Description
Jun 04, 2024 18:17:34 GMT+08:00	info	[DRS checked] lines compare complete
Jun 04, 2024 18:17:33 GMT+08:00	info	[DRS checked] lines compare start
Jun 04, 2024 18:17:16 GMT+08:00	info	[DRS checked] objects compare complete
Jun 04, 2024 18:17:15 GMT+08:00	info	[DRS checked] objects compare start
Jun 04, 2024 18:16:49 GMT+08:00	info	Check why a subtask is stopped. A subtask will be automatically stopped and deleted after the setting is complete.
Jun 04, 2024 18:16:48 GMT+08:00	info	release child transfer completed
Jun 04, 2024 18:16:48 GMT+08:00	info	drds transfer completed
Jun 04, 2024 18:16:45 GMT+08:00	info	drds [DRMS_FUNC] transfer completed
Jun 04, 2024 18:16:45 GMT+08:00	info	set the new full transfer sub-point:mysql-bin-000244101:17167740-2215-114532b4-1010a028117-1-624820
Jun 04, 2024 18:16:21 GMT+08:00	info	drds [DRMS_FUNC] transfer start

In addition, DRS can interconnect with Log Tank Service (LTS). After you enable log reporting to LTS, all logs generated by DRS instances will be uploaded to LTS for management. For details, see [Log Reporting](#).

----End

7.4 Data Comparison (Comparing Synchronization Items)

Scenarios

This section describes how to compare synchronization items to check if there are any differences between source and destination databases. To minimize the impact on services and shorten the service interruption duration, the following comparison methods are provided:

- Object-level comparison: compares objects such as databases, indexes, tables, views, stored procedures, functions, and table sorting rules.
- Data-level comparison is classified into row comparison and value comparison.

- Row comparison: It helps you compare the number of rows in the tables to be synchronized. This comparison method is recommended because it is fast.
- Value comparison: It helps you check whether data in the synchronized table is consistent. The comparison process is relatively slow.
- Sampling comparison: If there is a lot of data, it takes a long time to compare the number of rows and values. Sampling comparison is recommended. You can set an appropriate sampling ratio to complete a comparison faster.
- Account comparison: It compares usernames and permissions of the source and destination databases.
- Periodic comparison: DRS periodically compares the number of rows in the source database table with those in the destination database table and displays the comparison results. To compare objects periodically, enable [comparison policy](#).

When you check data consistency, compare the number of rows first. If the number of rows are inconsistent, you can then compare the data in the table to determine the inconsistent data.

Table 7-10 Supported comparison methods

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
To the clou d	MySQL->MySQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	MySQL->PostgreSQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	MySQL -> GaussDB distributed	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	MySQL -> GaussDB primary/standby	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
To the clou d	MySQL->GaussDB(for MySQL)	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	MySQL->GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	MySQL->MariaDB	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	PostgreSQL->PostgreSQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Suppo rted
To the clou d	PostgreSQL- >GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	PostgreSQL -> GaussDB primary/standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	PostgreSQL -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	DDM->MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	DDM->GaussDB(DWS)	Not suppor ted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
To the clou d	DDM->DDM	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Oracle->MySQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	Oracle->GaussDB(for MySQL)	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	Oracle -> GaussDB primary/standby	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	Oracle -> GaussDB distributed	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	Oracle->DDM	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Oracle->GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Oracle->PostgreSQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	DB2 for LUW -> GaussDB primary/standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
To the clou d	DB2 for LUW -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	DB2 for LUW->GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	TiDB->GaussDB(for MySQL)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server->GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server -> GaussDB primary/standby	Suppo rted	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server -> GaussDB distributed	Suppo rted	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server->Microsoft SQL Server	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server->MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Microsoft SQL Server->GaussDB(for MySQL)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
To the clou d	Microsoft SQL Server->PostgreSQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	MongoDB->DDS	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	MariaDB->MariaDB	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	MariaDB->MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	MariaDB->GaussDB(for MySQL)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	GaussDB(for MySQL)- >GaussDB(for MySQL)	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
To the clou d	Cassandra->GeminiDB Cassandra	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
To the clou d	Dynamo->GeminiDB Dynamo	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	MySQL->MySQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
From the clou d	MySQL->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	MySQL->CSS/ES	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	MySQL->Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	MySQL->MariaDB	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	DDM->MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	DDM->Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	DDM->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	DDS->MongoDB	Suppo rted	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted
From the clou d	DDS->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
From the clou d	PostgreSQL->PostgreSQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
From the clou d	PostgreSQL->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > GaussDB(DWS)	Suppo rted	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB primary/standby - > GaussDB primary/ standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB distributed -> MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
From the clou d	GaussDB distributed -> Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB distributed -> GaussDB(DWS)	Suppo rted	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted
From the clou d	GaussDB distributed -> Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB distributed -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB distributed- >GaussDB primary/standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB(for MySQL)- >MySQL	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
From the clou d	GaussDB(for MySQL)- >GaussDB(DWS)	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB(for MySQL)- >Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	GaussDB(for MySQL)- >CSS/ES	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
From the clou d	GaussDB(for MySQL)- >Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
From the clou d	MariaDB->MariaDB	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
From the clou d	Microsoft SQL Server- >Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	Oracle->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	Oracle -> GaussDB primary/standby	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
Self- built -> Self- built	Oracle -> GaussDB distributed	Suppo rted	Suppo rted	Suppo rted	Suppo rted	Not suppor ted
Self- built -> Self- built	MySQL->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	MySQL->CSS/ES	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync hron izati on Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
Self- built -> Self- built	MySQL -> GaussDB primary/standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	MySQL -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	PostgreSQL->Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB primary/standby - > MySQL	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB primary/standby - > Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB primary/standby - > Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB primary/standby - > GaussDB primary/ standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB distributed -> Oracle	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted

Sync roniza tion Dire ctio n	Data Flow	Objec t-level Comp arison	Row Comp arison	Value Comp arison	Dyna mic Comp arison	Accou nt- level Comp arison
Self- built -> Self- built	GaussDB distributed -> Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	GaussDB distributed -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	DB2 for LUW -> GaussDB primary/standby	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	DB2 for LUW -> GaussDB distributed	Suppo rted	Suppo rted	Not suppor ted	Not suppor ted	Not suppor ted
Self- built -> Self- built	Microsoft SQL Server- >Kafka	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted	Not suppor ted

Constraints

- You can manually create a comparison task only when the task is in the incremental phase.
- During a comparison, the comparison items are case sensitive. If one of the source or destination database is case insensitive and the other one is case sensitive, the comparison result may be inconsistent.
- When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. If operations are performed on data in the source database during data comparison, the comparison results may be inconsistent.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.

- If data in the destination database is modified separately, the comparison results may be inconsistent.
- If the encoding of the source database character type is abnormal, the database driver will convert the character type to an abnormal code point during DRS migration or comparison. As a result, the values may be consistent but the bytes may be inconsistent.
- Currently, only tables with primary keys support value comparison. For tables that do not support value comparison, you can compare rows. Therefore, you can compare data by row or value based on scenarios.
- The DRS task cannot be suspended during value comparison. Otherwise, the comparison task may fail.
- Some data types do not support value comparison. For details, see [Which of the Following Data Types Are Not Supported By Value Comparison?](#)
- To prevent resources from being occupied for a long time, DRS limits the comparison duration. If the comparison duration exceeds the threshold, the comparison task stops automatically.
 - When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. The comparison duration limits to 30 minutes. After the threshold, the comparison tasks automatically stop and the full migration task stops.
 - For a row comparison task manually created in the incremental phase, if the source database is a relational database, the row comparison duration limits to 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration limits to 30 minutes.
- To avoid occupying resources, the comparison results of DRS tasks can be retained for a maximum of 60 days. After 60 days, the comparison results are automatically cleared.
- For a migration task from MySQL, virtual columns in the source database do not support value comparison. During the comparison, virtual columns are filtered out.
- If you want to perform value comparison or sampling comparison for a DRS task, select large or higher specifications when creating the DRS task.
- In the many-to-one row comparison scenario, the number of rows in the table in the source database is compared with that in the aggregation table mapped to the destination database.
- In the many-to-one synchronization scenario, value comparison is not recommended because data consistency cannot be ensured.
- Value comparison is not supported for a task in which tables in one database are mapped to multiple databases.
- For a synchronization task from MySQL, GaussDB(for MySQL), or MariaDB, virtual columns in the source database do not support value comparison. During the comparison, virtual columns are filtered out.
- If the source is a PostgreSQL database, the index and constraint names will be changed during table mapping. As a result, the index and constraint names are inconsistent.
- If a table in the source MySQL database contains a binary field with a fixed length, the MySQL driver adds \0 to the end of the data based on the length.

As a result, there may be data inconsistency after the data is synchronized to the destination GaussDB database.

- The empty character inserted into an Oracle database is processed as NULL. For tasks whose destination is an Oracle database, an empty string is considered as NULL. If data in the source database is empty and that in the destination database is NULL, the comparison result is consistent.
- During value comparison for synchronization from Oracle to GaussDB Distributed, if the LOB comparison policy is set to **Compare length**, the BLOB comparison is ignored because BLOB data in the distributed GaussDB instance fails to be queried using **DBE_LOB.LOB_GET_LENGTH**.
- When data is filtered during value comparison from Dynamo to GeminiDB Dynamo, the filter criteria do not support the binary type. Value comparison is based on the source database data. The scenario where the destination database has more data records than the source database cannot be identified.

Impact on Databases

- Object comparison: System tables of the source and destination databases are queried, occupying about 10 sessions. The database is not affected. However, if there are a large number of objects (for example, hundreds of thousands of tables), the database may be overloaded.
- Row comparison: The number of rows in the source and destination databases is queried, which occupies about 10 sessions. The SELECT COUNT statement does not affect the database. However, if a table contains a large amount of data (hundreds of millions of records), the database will be overloaded and the query results will be returned slowly.
- Value comparison: All data in the source and destination databases is queried, and each field is compared. The query pressure on the database leads to high I/O. The query speed is limited by the I/O and network bandwidth of the source and destination databases. Value comparison occupies one or two CPUs, and about 10 sessions.
- Account comparison: The accounts and permissions of the source and destination databases are queried, which does not affect the database.

Estimated Comparison Duration

- Object comparison: Generally, the comparison results are returned within several minutes based on the query performance of the source database. If the amount of data is large, the comparison may take dozens of minutes.
- Row comparison: The SELECT COUNT method is used. The query speed depends on the database performance.
- Value comparison: If the database workload is not heavy and the network is normal, the comparison speed is about 5 MB/s.
- Account comparison: The results are returned with the object-level comparison results. If the number of objects is small, the results are returned in several minutes.

Prerequisites

- You have logged in to the DRS console.

- A synchronization task has been started.

Creating a Comparison Task

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 Choose **Synchronization Comparison**.

Step 3 Compare synchronization items.


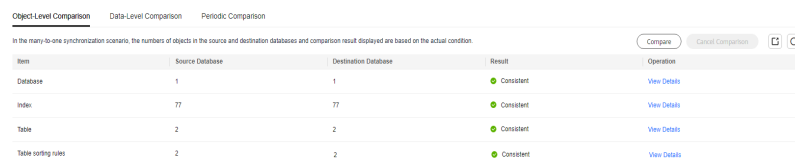
- Create an **Object-Level Comparison** task. On the **Object-Level Comparison** tab, click **Compare**. Wait for a while and click  to check whether the comparison results of the source and destination databases are consistent. Locate a comparison item you want to view and click **View Details** in the **Operation** column.

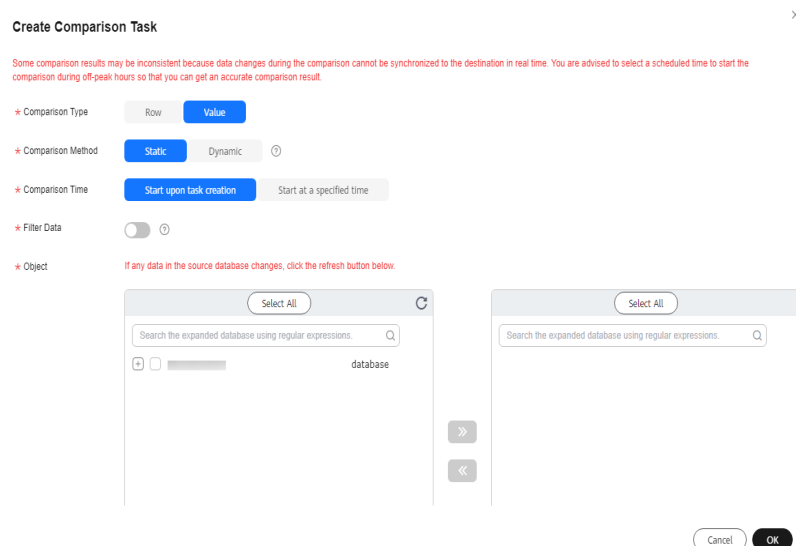
Figure 7-20 Object-level comparison



Item	Source Database	Destination Database	Result	Operation
Database	1	1	Consistent	View Details
Index	77	77	Consistent	View Details
Table	2	2	Consistent	View Details
Table sorting rules	2	2	Consistent	View Details

- On the **Data-Level Comparison** (row comparison and value comparison) tab, click **Create Comparison Task**. In the displayed dialog box, specify **Comparison Type**, **Comparison Method**, **Comparison Time**, and **Object**. Then, click **OK**.

Figure 7-21 Creating a data-level comparison task



Create Comparison Task

Some comparison results may be inconsistent because data changes during the comparison cannot be synchronized to the destination in real time. You are advised to select a scheduled time to start the comparison during off-peak hours so that you can get an accurate comparison result.

- * Comparison Type: Row Value
- * Comparison Method: Static Dynamic
- * Comparison Time: Start upon task creation Start at a specified time
- * Filter Data:
- * Object: [Refresh](#)
If any data in the source database changes, click the refresh button below.

Search the expanded database using regular expressions. Q

database

Cancel OK

- **Comparison Type:** compares rows and values.
 - Row comparison: checks whether the source table has the same number of rows as the destination table.

 NOTE

- After a task enters the incremental comparison phase, you can create a row comparison task.
- For Oracle to GaussDB synchronization, the row comparison task is automatically triggered after the full synchronization is complete.
- **Value comparison:** checks whether the source table has the same data as the destination table.

 NOTE

- After a task enters the incremental synchronization phase, you can create a value comparison task. After the full synchronization is complete, data in the source database cannot be changed. Otherwise, the comparison result will be inconsistent.
- After the synchronization from Oracle to GaussDB primary/standby and from Oracle to GaussDB distributed enters the incremental verification phase, the comparison service starts to extract data from the incremental logs of the tables involved in the task for continuous comparison. If you want to compare the status of other tables, you can cancel the ongoing comparison task and create a comparison task.

Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.

- **Comparison Policy:** DRS supports one-to-one and many-to-one comparison policies.
 - **One-to-one:** compares the number of rows in a table in the source database with that in the table mapped to the destination database.
 - **Many-to-one:** compares the number of rows in a table in the source database with that in the aggregate table mapped to the destination database.

 NOTE

If you select **Row Comparison** for **Comparison Type**, the **Comparison Policy** option becomes available.

- **Comparison Method:** DRS provides static and dynamic comparison methods.
 - **Static:** All data in the source and destination databases is compared. The comparison task ends as the comparison is completed. Static comparison can only be performed when there are no ongoing services.
 - **Dynamic:** All data in the source database is compared with that in the destination database. After the comparison task is complete, incremental data in the source and destination databases is compared in real time. A dynamic comparison can be performed when data is changing.

 NOTE

- If you select **Value** for **Comparison Type**, the **Comparison Method** option becomes available.
- During database-level synchronization, tables cannot be created in the source database during dynamic comparison. If you want to create a table in the source database, cancel the dynamic comparison first. After the new table is created and synchronized, restart the dynamic comparison.
- **Comparison Time:** You can select **Start upon task creation** or **Start at a specified time**. There is a slight difference in time between the source and destination databases during synchronization. Data inconsistency may occur. You are advised to compare migration items during off-peak hours for more accurate results.
- **LOB Comparison Policy:** The value can be **Ignore LOB comparison**, **Compare length**, **Compare hash values**, or **Compare content**.

 NOTE

LOB comparison policy can be set only for data synchronization from Oracle to GaussDB.

- **Ignore LOB comparison:** The system ignores LOB data during value comparison. You are advised to select **Ignore LOB comparison** because comparing LOB data increases the database load, depending on the LOB comparison method and data volume. Evaluate and test the LOB comparison policy based on the source and destination databases to ensure database performance and stability.
- **Compare length:** The built-in functions of the source and destination databases are used to obtain the LOB data length for data comparison.
- **Compare hash values:** The built-in functions of the source and destination databases are used to obtain the LOB data hash values for data comparison. Oracle databases use the HASH function in the DBMS_CRYPTO package to obtain the LOB data hash values. To use the DBMS_CRYPTO package, grant SYSDBA permissions to the user.
Reference statement:
GRANT EXECUTE ON DBMS_CRYPTO TO USER;
- **Compare content:** The source database reads data in streaming mode and then performs hashing. The destination database uses built-in functions to obtain the LOB data hash values. Compared with hash value comparison, this method reduces the pressure on the source database, but it takes longer.
- **Filter Data:** After this function is enabled, objects can be compared based on the configured filtering criteria.

 NOTE

Data filtering and comparison can be set only for synchronization tasks from Oracle to GaussDB, GaussDB to Oracle, GaussDB to GaussDB, MySQL to MySQL and MySQL to GaussDB.

- i. After enabling **Filter Data**, add filtering criteria for the table objects to be compared.

- ii. In the **Filtering Criteria** area, enter the filtering criteria, and click **Verify**.

 **NOTE**

- Each table has only one verification rule.
 - Up to 512 tables can be filtered at a time. If there are more than 512 tables, perform rule verifications in batches.
 - Standard SQL statements can be used to filter records. Each expression cannot contain packages, functions, variables, or constants specific to a database engine.
 - Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, sid > 3 and sname like "G %". A maximum of 512 characters are allowed.
 - In SQL statements for setting filter criteria, keywords must be marked with a field identifier, and the values of **datetime** (including date and time) and character string type must be enclosed in single quotation marks, for example, ``update` > '2022-07-13 00:00:00'` and `age >10, `update` ='abc'`.
 - If the **TIMESTAMP** type is used as a filtering condition, the time of the character type must be set to the time value in the UTC time zone. For example, in MySQL, the **TIMESTAMP** data is stored based on the UTC time zone. You need to use the time value in the UTC time zone for comparison.
 - Implicit conversion rules are not supported. Enter filtering criteria of a valid data type. For example, if column c of an Oracle database uses characters of the **varchar2** type, the filtering criteria must be set to `c > '10'` instead of `c > 10`.
 - Filter criteria cannot be configured for large objects, such as **CLOB**, **BLOB**, and **BYTEA**.
 - Non-idempotent expressions or functions cannot be used as data processing conditions, such as **SYSTIMESTAMP** and **SYSDATE**, because the returned result may be different each time the function is called.
 - During data filtering for real-time synchronization with Oracle serving as the source database, the fixed-length character types **NCHAR** and **CHAR** must be matched using complete fixed-length characters.
 - You are not advised to set filter criteria for fields of approximate numeric types, such as **FLOAT**, **DECIMAL**, and **DOUBLE**.
 - Do not use fields containing special characters as a filter condition.
 - Objects whose database names, schema names, or table names are case insensitive cannot be filtered and compared.
 - Currently, condition-based filtering is not supported when there are more than 50,000 tables in a database.
 - For security purposes, keywords or functions with update meanings, such as the **for update** statement and **updatexml** function, cannot be used in SQL fragments.
- iii. After the verification is successful, click **Generate Processing Rule**. The rule is displayed.
 - iv. Click **OK**.
 - **Object**: You can select objects to be compared based on the scenarios.
 - Create a **Data-Level Comparison (sampling comparison)** task. If there is a lot of data, it takes a long time to compare the number of rows and values.

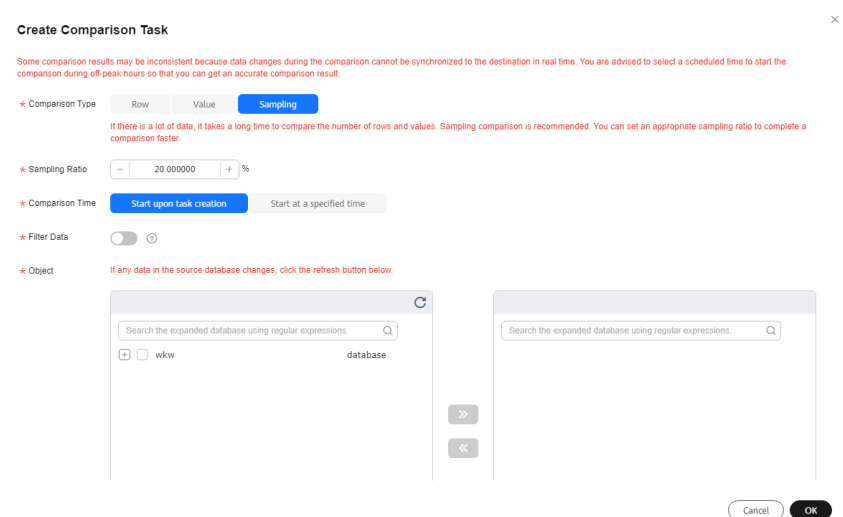
Sampling comparison is recommended. You can set an appropriate sampling ratio to complete a comparison faster.

On the **Data-Level Comparison** page, click **Create Comparison Task**, select **Sample** for **Comparison Type**, set **Sampling Ratio**, **Comparison Time**, and **Object**, and click **OK**.

 **NOTE**

- Only real-time synchronization tasks from MongoDB to DDS, DDS to MongoDB, and GaussDB Distributed to GaussDB Distributed support sampling comparison.
- Only tasks with large or higher specifications support sampling comparison.
- The **\$sample** command provided by MongoDB is used for data sampling: **{ \$sample: { size: <positive integer N> } }**.
In the preceding command, **size** is fixed to 1000. To perform MongoDB sampling comparison provided by DRS, the **\$sample** command is executed for multiple times until the sampled data volume reaches the specified sampling ratio.
- In a single sampling comparison task, the **\$sample** command is executed for multiple times. In this case, the same data in the source database may be obtained for multiple times.
- The total data volume of a collection is obtained using **estimatedDocumentCount**. It is an estimated value instead of an accurate value. The actual sampling data volume and sampling ratio may fluctuate, which is normal.
- Sampling comparison does not support documents whose **_id** is of the BinData data type. Documents whose **_id** is of the BinData data type will be filtered out and not be compared.
- If the number of inconsistent data records in a single MongoDB sampling comparison task exceeds 10,000, the comparison stops. A maximum of 10,000 inconsistent data records can be displayed.

Figure 7-22 Sampling comparison



- Create an **Account-Level Comparison** task. Click the **Account-Level Comparison** tab to view the comparison results of database accounts and permissions.

 **NOTE**

- Full synchronization tasks do not support account comparisons.
- Only PostgreSQL to PostgreSQL synchronization supports account comparison.


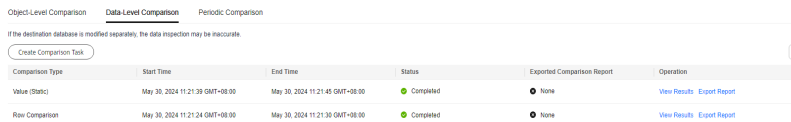
Step 4 After the comparison creation task is submitted, the **Data-Level Comparison** tab is displayed. Click  to refresh the list and view the comparison result of the specified comparison type.

Figure 7-23 Data-level comparison



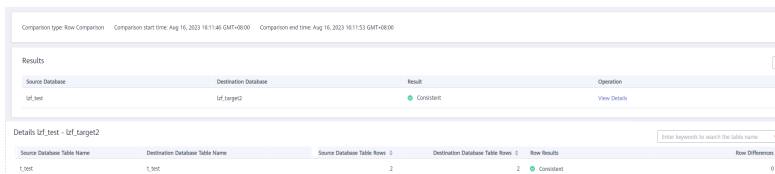
Comparison Type	Start Time	End Time	Status	Exported Comparison Report	Operation
Value (Static)	May 30, 2024 11:21:39 GMT+08:00	May 30, 2024 11:21:45 GMT+08:00	Completed	None	View Results Export Report
Row Comparison	May 30, 2024 11:21:24 GMT+08:00	May 30, 2024 11:21:30 GMT+08:00	Completed	None	View Results Export Report

Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.

If you want to view the row or value comparison details, click **View Results**.

If you want to download the row comparison or value comparison result, locate a specified comparison type and click **Export Report** in the **Operation** column.

Figure 7-24 Row comparison details



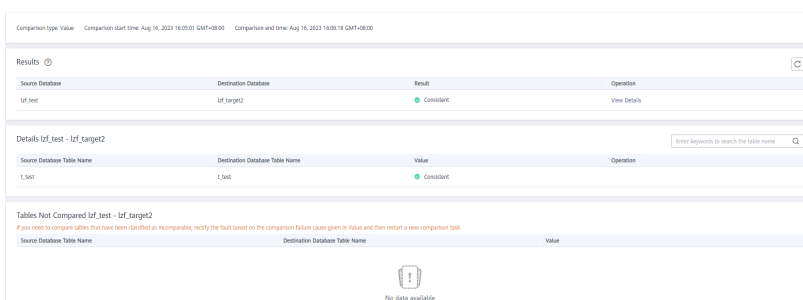
Comparison type: Row Comparison Comparison start time: Aug 16, 2023 16:11:46 GMT+08:00 Comparison end time: Aug 16, 2023 16:11:53 GMT+08:00

Source Database	Destination Database	Result	Operation
lf_test	lf_target2	Consistent	View Details

Details lf_test - lf_target2

Source Database Table Name	Destination Database Table Name	Source Database Table Rows	Destination Database Table Rows	Row Results	Row Differences
t_test	t_test	2	2	Consistent	0

Figure 7-25 Value comparison details



Comparison type: Value Comparison start time: Aug 16, 2023 16:05:01 GMT+08:00 Comparison end time: Aug 16, 2023 16:06:18 GMT+08:00


Source Database	Destination Database	Result	Operation
lf_test	lf_target2	Consistent	View Details

Details lf_test - lf_target2

Source Database Table Name	Destination Database Table Name	Value	Operation
t_test	t_test	Consistent	

Tables Not Compared lf_test - lf_target2

If you need to compare tables that have been classified as incompatible, rectify the fault based on the comparison failure cause given in Value and then restart a new comparison task.

Source Database Table Name	Destination Database Table Name	Value
 No data available		

 **NOTE**

- You can cancel a running task at any time and view the comparison report of a canceled comparison task.
- You can sort the row comparison results displayed on the current page in ascending or descending order based on the number of rows in the source database table or the destination database table.
- If a negative number is displayed in the **differences** column, the number of rows in the destination database table is greater than that in the source database table. If a positive number is displayed in the **differences** column, the number of rows in the source database table is greater than that in the destination database table.

----End

Periodic Comparison

Periodic comparison indicates that DRS periodically compares the number of rows in the source database table with those in the destination database table and displays the comparison results.

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 Click the **Synchronization Comparison** tab.

Step 3 Click the **Periodic Comparison** tab and click **Modify Comparison Policy** to modify the comparison policy.

 **NOTE**

Only MySQL-to-MySQL, MySQL-to-GaussDB(for MySQL), MySQL-to-GaussDB(DWS), and GaussDB(for MySQL)-to-MySQL synchronization tasks support periodic comparison.

Step 4 In the **Modify Comparison Policy** dialog box, enable periodic comparison, configure the comparison frequency and time, and click **Yes**.

 **NOTE**

- After periodic comparison is enabled, DRS compares the number of rows at the scheduled time. You can view the comparison results on the **Data-Level Comparison** tab.
- After periodic comparison is disabled, only historical comparison results can be viewed.
- Modifications to the comparison policy settings take effect from the next comparison and do not affect the on-going periodic comparison tasks.
- During periodic comparison, the source and destination databases will be read. Perform the comparison during off-peak hours.
- During periodic comparison, ultra-large tables (those with more than 100 million rows) are automatically filtered out. You can use data-level comparison to spot check such large tables. It is not recommended that these large tables be compared periodically.

Figure 7-26 Modify Comparison Policy

----End

7.5 Managing Objects

7.5.1 Editing Synchronization Objects

This section describes how to edit synchronization objects in an incremental synchronization task.

- For a normal incremental task, you can edit synchronization objects by adding or deleting databases and tables to be synchronized. DRS will synchronize the selected objects in newly added tables.
- For a failed incremental task, you can edit synchronization objects by changing the objects to be synchronized. If an incremental synchronization fails due to incorrect synchronization objects, you can use this function to remove the databases or tables and submit the task again to restore the task.

- During synchronization object editing of a failed incremental task, do not resume or reset the synchronization task before it is started. Otherwise, the synchronization object cache will be cleared, and the synchronization objects fail to be edited.
- If an incremental task fails to be started for the first time and the data capture and parsing fail, ensure that the capture and parsing processes of the primary task are normal, and then edit the task to add or delete tables.
- When you edit synchronization objects to add or delete tables, if a subtask fails due to an exception after the task is started, information about the added or deleted tables will not be saved in the synchronization objects after the subtask is stopped. In addition, data in these tables is synchronized from the source database to the destination database. As a result, data in the destination database may be inconsistent with that in the source database.
- During the editing, you cannot pause, reset, or restart a parent task, change specifications, or reset the start position. You cannot reset the password of the source or destination database. You cannot modify the database connection information (IP address and password). Otherwise, data may be inconsistent or the task may fail. You can pause, resume, stop a subtask.

For details about the synchronization scenarios where synchronization objects can be edited in the incremental phase, see [Real-time Synchronization Scenarios Where Synchronization Objects Can Be Edited](#).


Prerequisites

You have logged in to the DRS console.

Method 1

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task and click **Edit** in the **Operation** column.

Figure 7-27 Editing synchronization objects



- Step 2** On the **Set Synchronization Task** page, change the objects to be synchronized and click **Next**.

- You can search the expanded database using regular expressions.
- If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed.
- The name of the selected synchronization object cannot contain spaces.

- Step 3** On the **Process Data** page, set rules for a new table by referring to [Processing Data](#).

NOTE

The processing rules for a synchronized table cannot be modified.

- Step 4** On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 5 On the **Confirm Task** page, specify **Start Time**, confirm that the configured information is correct, and click **Submit**.

Step 6 Go back to the **Data Synchronization Management** page. In the synchronization task list, the current task status is **Incremental synchronization**, and a subtask in the **Modifying task** status is generated. You can click the subtask name and click the **Synchronization Progress** tab to view the synchronization progress of each synchronization object. The streaming progress display helps you learn about the synchronization progress in real time.

Step 7 After the subtask change is complete, incremental synchronization is performed for the edited synchronization objects.

----End

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task.

Step 2 On the displayed page, click the **Synchronization Mapping** tab and click **Edit** to the right of the synchronization object.

Step 3 Perform [Step 2](#) to [7](#) from method 1.

----End

Real-time Synchronization Scenarios Where Synchronization Objects Can Be Edited

To the cloud

- MySQL -> MySQL
- MySQL -> GaussDB(for MySQL)
- MySQL -> GaussDB Primary/Standby
- MySQL -> GaussDB Distributed
- MySQL->MariaDB
- DDM -> DDM
- DDM -> GaussDB(DWS)
- Oracle -> MySQL
- Oracle -> PostgreSQL

- Oracle -> GaussDB Primary/Standby
- Oracle -> GaussDB Distributed
- Oracle -> GaussDB(DWS)
- Oracle -> GaussDB(for MySQL)
- PostgreSQL -> PostgreSQL
- PostgreSQL -> GaussDB Distributed
- PostgreSQL -> GaussDB Primary/Standby
- MariaDB -> MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- DB2 for LUW -> GaussDB Distributed
- DB2 for LUW -> GaussDB Primary/Standby
- DB2 for LUW -> GaussDB(DWS)
- Microsoft SQL Server->MySQL
- Microsoft SQL Server->GaussDB(for MySQL)
- Microsoft SQL Server -> GaussDB(DWS)
- Microsoft SQL Server -> GaussDB Distributed
- Microsoft SQL Server -> GaussDB Primary/Standby
- Microsoft SQL Server->PostgreSQL
- GaussDB(for MySQL) -> GaussDB(for MySQL)

From the cloud

- MySQL -> MySQL
- MySQL -> CSS/ES
- MySQL -> Kafka
- MySQL -> Oracle
- MySQL->MariaDB
- DDM -> Kafka
- DDS -> Kafka
- PostgreSQL -> PostgreSQL
- PostgreSQL -> Kafka
- MariaDB -> MariaDB
- GaussDB(for MySQL) -> MySQL
- GaussDB(for MySQL) -> CSS/ES
- GaussDB(for MySQL) -> GaussDB(DWS)
- GaussDB(for MySQL) -> Oracle
- GaussDB(for MySQL) -> Kafka
- GaussDB Primary/Standby -> MySQL
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> Kafka
- GaussDB Primary/Standby -> GaussDB(DWS)

- GaussDB Primary/Standby -> GaussDB Distributed
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Distributed -> MySQL
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> Kafka
- GaussDB Distributed -> GaussDB(DWS)
- GaussDB Distributed -> GaussDB Distributed
- GaussDB Distributed -> GaussDB Primary/Standby
- Microsoft SQL Server->Kafka

Self-built -> Self-built

- MySQL -> CSS/ES
- MySQL -> Kafka
- MySQL -> GaussDB Primary/Standby
- Oracle -> Kafka
- Oracle -> GaussDB Primary/Standby
- Oracle -> GaussDB Distributed
- PostgreSQL -> Kafka
- DB2 for LUW -> GaussDB Distributed
- DB2 for LUW -> GaussDB Primary/Standby
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> Kafka
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> Kafka
- GaussDB Distributed -> GaussDB Distributed
- Microsoft SQL Server->Kafka

7.5.2 Importing Synchronization Objects

Real-time synchronization supports the import of objects through files. After a task is created, you can import object files on the **Set Synchronization Task** page.

Precautions


- Only Windows Microsoft Excel 97-2003 (*.xls), 2007, and later (*.xlsx) files can be imported. The downloaded compressed package provides the templates of the two versions.
- The file name can contain only spaces, letters, digits, hyphens (-), underscores (_), and parentheses.
- The object information in the template must be filled in according to the format in the Excel file. The database object name (such as the database name, schema name, and table name) is case sensitive cannot contain the following characters: <>.". Object names starting or ending with spaces are not supported.

- The task in the configuration supports table-level synchronization, database-level synchronization, or file import mode. Each time you switch to a new mode, the selected or imported database objects are cleared, and you need to select or import them again.
- If you want to import a file for mapping, fill in the first and second columns of the file based on the template. If the first two columns of a row are left blank, the row will be ignored.
- For the task created using the file import mode, database-level and table-level synchronization are not supported after the task is started.
- If you edit a task, the imported file must contain information about all objects. Importing only the updated objects is not allowed.
- If you edit a task again, the objects that have been synchronized cannot be mapped again. Ensure that the object names remain unchanged after the mapping.
- If you edit a task again, the exported object information is the synchronized object information.
- If the verification fails after the file is uploaded, click **View Failure Details** to download the error information.
- The object names entered in the Excel file must use the same letter case as the source object names.

Procedure

Step 1 On the **Set Synchronization Task** page, click **Import object file** in the **Synchronization Object** field.

Figure 7-28 Importing an object file



Step 2 Click **Download Template**.

Step 3 Download the template and enter information about the objects to be imported.

Step 4 Click **Select File**. In the displayed dialog box, select the edited template.

Step 5 Click **Upload**.

----End

7.5.3 Changing Object Names (Mapping Object Names)

Data synchronization allows you to synchronize objects (including databases, schemas and tables) in a sources database to the corresponding objects in a destination database. Object name mapping can be used only in the following scenarios:

- **Changing an Object Name:** If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one.
- **Many-to-One Synchronization:** To synchronize databases, schemas, and tables in many-to-one scenarios, you can use object name mapping.

This section describes how to map objects when configuring a data synchronization task. For details about the mapping relationship, see [Viewing Synchronization Mapping Information](#).

Precautions

- Objects whose database names or table names contain newline characters cannot be mapped.
- Object name mapping can be used for the first time you select synchronization objects for a data synchronization task.
- Object name mapping can be used for the first time you add or delete the synchronization object which is not in a mapping relationship.
- If the destination DB is a type of PostgreSQL (for example, GaussDB(DWS), GaussDB, or PostgreSQL), the destination schema name cannot start with **pg_**. Otherwise, the migration fails.
- The following data flow scenarios support basic mapping:
 - AWS DocumentDB -> DDS
 - Cassandra -> GeminiDB Cassandra
 - MySQL -> MariaDB
 - MySQL -> GaussDB Primary/Standby
 - MySQL -> GaussDB Distributed
 - MySQL -> GaussDB(DWS)
 - MySQL -> GaussDB(for MySQL)
 - MySQL -> Kafka
 - MySQL -> PostgreSQL
 - MySQL -> Oracle
 - MariaDB -> MariaDB
 - MariaDB -> MySQL
 - MariaDB -> GaussDB(for MySQL)
 - DDM -> DDM
 - DDM -> Oracle
 - DDM -> GaussDB(DWS)
 - DB2 for LUW -> GaussDB Primary/Standby
 - DB2 for LUW -> GaussDB Distributed
 - DB2 for LUW -> GaussDB(DWS)
 - DDS -> MongoDB
 - Microsoft SQL Server -> GaussDB(DWS)
 - Microsoft SQL Server -> GaussDB Primary/Standby
 - Microsoft SQL Server -> GaussDB Distributed

- Microsoft SQL Server -> Microsoft SQL Server
- Microsoft SQL Server -> MySQL
- Microsoft SQL Server -> PostgreSQL
- Microsoft SQL Server -> GaussDB(for MySQL)
- MongoDB -> DDS
- GaussDB Primary/Standby -> MySQL
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Primary/Standby -> GaussDB Distributed
- GaussDB Distributed -> MySQL
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> GaussDB Distributed
- GaussDB Distributed -> GaussDB Primary/Standby
- GaussDB(for MySQL) -> MySQL
- GaussDB(for MySQL) -> GaussDB(DWS)
- GaussDB(for MySQL) -> GaussDB(for MySQL)
- Oracle -> MySQL
- Oracle -> GaussDB(for MySQL)
- Oracle -> GaussDB(DWS)
- Oracle -> GaussDB Primary/Standby
- Oracle -> GaussDB Distributed
- Oracle -> PostgreSQL
- Oracle -> DDM
- PostgreSQL -> PostgreSQL
- PostgreSQL -> GaussDB(DWS)
- PostgreSQL -> GaussDB Primary/Standby
- PostgreSQL -> GaussDB Distributed
- TiDB -> GaussDB(for MySQL)
- The following data flow scenarios support only table-level mapping:
 - Dynamo -> GeminiDB Dynamo
 - MySQL -> CSS/ES
 - GaussDB(for MySQL) -> CSS/ES
- The following data flow scenarios support only database-level mapping:
 - GaussDB Primary/Standby -> GaussDB(DWS)
 - GaussDB Distributed -> GaussDB(DWS)
- The following data flow types do not support schema mapping for tables of the smallserial, serial, and bigserial types:
 - GaussDB Primary/Standby -> GaussDB Distributed
 - GaussDB Primary/Standby -> GaussDB Primary/Standby
 - GaussDB Distributed -> GaussDB Distributed

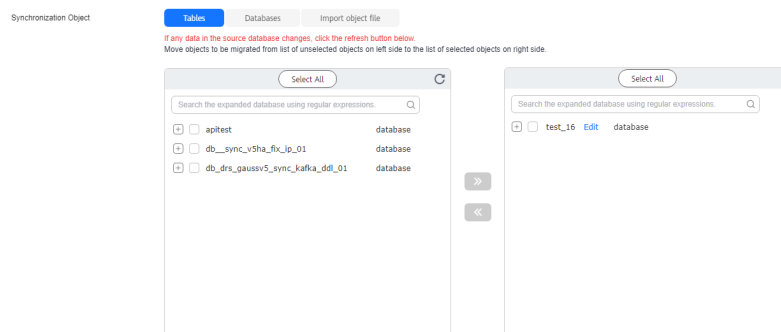
- GaussDB Distributed -> GaussDB Primary/Standby
- The following data flow types support many-to-one object mapping:
 - MySQL -> MySQL
 - MySQL -> GaussDB(for MySQL)
 - MySQL -> GaussDB(DWS)
 - MySQL -> Kafka
 - MySQL -> Oracle
 - GaussDB(for MySQL) -> MySQL
 - GaussDB(for MySQL) -> GaussDB(for MySQL)
 - GaussDB(for MySQL) -> GaussDB(DWS)
 - Oracle -> GaussDB
 - Oracle -> GaussDB(DWS)
 - GaussDB Primary/Standby -> GaussDB(DWS)
 - GaussDB Distributed -> GaussDB(DWS)
 - MySQL->Oracle
 - GaussDB Primary/Standby -> GaussDB Distributed
 - GaussDB Primary/Standby -> GaussDB Primary/Standby
 - GaussDB Distributed -> GaussDB Distributed
 - GaussDB Distributed -> GaussDB Primary/Standby

Changing an Object Name

During real-time synchronization, if the names of source databases, schemas, or tables to be synchronized are different from those in the destination, you can map the source names to the destination ones. For example, when synchronizing database A in the source database to database B in the destination database, you need to map database name first.

- Step 1** On the **Set Synchronization Task** page, select the database that needs to be mapped from the synchronization objects on the right area and click **Edit**.

Figure 7-29 Mapping databases



- Step 2** Change a name.

In the displayed dialog box, enter a new object name. The new name is the name of the object saved in the destination database.

Figure 7-30 Changing a database name

Edit Database Name ×

i The new database name will be used in the destination database. ×

i After the database names are edited, refresh the pane on the right to obtain the latest status of the destination database objects. ×

New Database Name ?

Cancel OK

Figure 7-31 Changing a schema name

Edit Schema Name ×

i The new schema name will be used in the destination database. ×

New Schema Name ?

Cancel OK

Figure 7-32 Changing a table name

Edit Table Name ×

i The new table will be used in the destination database. ×

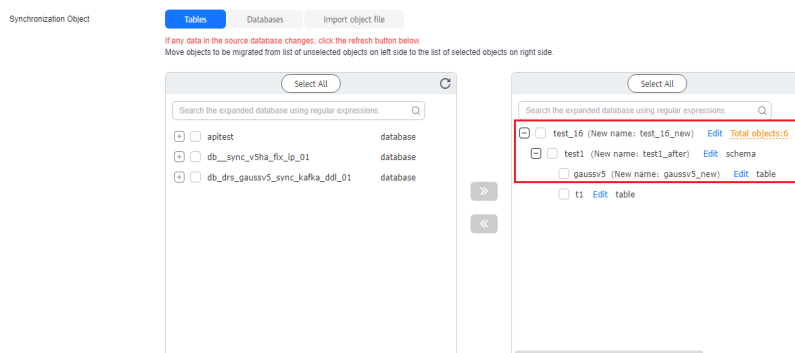
New Table Name ?

Cancel OK

Step 3 Check the result.

After the name is changed, the name before modification and the new name are displayed. The object name mapping is complete.

Figure 7-33 Checking the result



----End

Many-to-One Synchronization

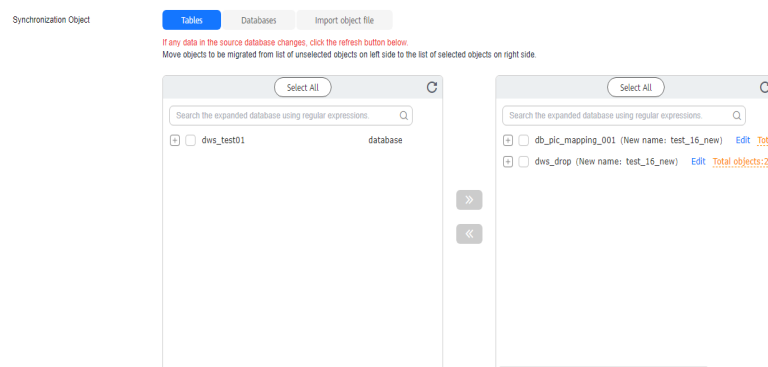
During real-time synchronization, you can use object name mapping to change the names of databases, schemas, and tables to enable many-to-one synchronization.

The following uses a database-level many-to-one synchronization as an example. For schema- and table-level many-to-one scenarios, perform the following similar steps:

- Step 1** On the **Set Synchronization Task** page, select the target databases that need to be mapped from the synchronization objects on the right area and click **Edit** one by one.
- Step 2** In the displayed dialog box, enter a given database name. The new database name is the name of the database saved in the destination DB instance.

For example, change the databases **test_16** and **auto_nprocess_001** in the source DB instance to the database **test_16_new**.

Figure 7-34 Checking the result



- Step 3** After the database names are changed, the database names before modification and the new database name are displayed. The many-to-one database mapping is complete.

Step 4 Wait until the task starts and the synchronization is complete. Then, the databases **test_16** and **auto_nprocess_001** can be synchronized to the database **test_16_new**.

----End

7.5.4 Viewing Synchronization Mapping Information

During real-time synchronization, the objects that can be mapped to the destination include databases, schemas, tables, and columns (in data processing). After a mapping relationship between objects is established, you can view details about the mapping.

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click the **Synchronization Mapping** tab to view the mapping details.

Figure 7-35 Viewing mapping information

Source Database	Table Name	Destination Database	Table Name	Operation
db001	L_mysql_2013_db001	db0012000	L_mysql_2013_db001	Full
db001	L_mysql_2013_db001_01	db0012000	L_mysql_2013_db001_012000	Full

NOTE

When you select an object, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed.

- Step 3** In the upper right corner, filter and search for the mapping relationships by object or column.

----End

7.5.5 Processing Data

DRS processes synchronized objects and allows you to add rules for selected objects. The processing rules supported by each data flow type are different. Currently, only some data flow types support data processing. For details, see [Table 7-11](#).

Table 7-11 Data flow types that support data processing

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Adition al Column	Column Processin g
To the cloud	MySQL->MySQL	Supporte d	Supporte d	Supporte d
To the cloud	MySQL -> GaussDB distributed	Supporte d	Supporte d	Supporte d
To the cloud	MySQL -> GaussDB primary/ standby	Supporte d	Supporte d	Supporte d
To the cloud	MySQL->GaussDB(DWS)	Supporte d	Supporte d	Not supported
To the cloud	MySQL->GaussDB(for MySQL)	Supporte d	Supporte d	Supporte d
To the cloud	MySQL->MariaDB	Supporte d	Supporte d	Supporte d
To the cloud	DDM->MySQL	Not supported	Not supported	Supporte d
To the cloud	DDM->GaussDB(DWS)	Not supported	Supporte d	Not supported
To the cloud	Oracle->GaussDB(DWS)	Supporte d	Supporte d	Not supported
To the cloud	Oracle->MySQL	Supporte d	Not supported	Not supported
To the cloud	Oracle->GaussDB(for MySQL)	Supporte d	Not supported	Not supported
To the cloud	Oracle -> GaussDB primary/ standby	Supporte d	Not supported	Supporte d

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Additional Column	Column Processin g
To the cloud	Oracle -> GaussDB distributed	Supporte d	Not supported	Supporte d
To the cloud	DB2 for LUW -> GaussDB primary/ standby	Supporte d	Not supported	Not supported
To the cloud	DB2 for LUW -> GaussDB distributed	Supporte d	Not supported	Not supported
To the cloud	MariaDB->MariaDB	Supporte d	Not supported	Not supported
To the cloud	MariaDB->MySQL	Supporte d	Supporte d	Supporte d
To the cloud	MariaDB->GaussDB(for MySQL)	Supporte d	Supporte d	Supporte d
From the cloud	MySQL->MySQL	Supporte d	Supporte d	Supporte d
From the cloud	MySQL->Kafka	Not supported	Not supported	Supporte d
From the cloud	MySQL->CSS/ES	Supporte d	Not supported	Supporte d
From the cloud	MySQL->Oracle	Supporte d	Not supported	Not supported
From the cloud	MySQL->MariaDB	Supporte d	Supporte d	Supporte d
From the cloud	DDM->MySQL	Not supported	Not supported	Supporte d

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Additional Column	Column Processin g
From the cloud	GaussDB primary/standby -> MySQL	Supporte d	Not supported	Not supported
From the cloud	GaussDB primary/standby -> Oracle	Supporte d	Not supported	Supporte d
From the cloud	GaussDB primary/standby -> Kafka	Not supported	Not supported	Supporte d
From the cloud	GaussDB primary/standby -> GaussDB(DWS)	Supporte d	Not supported	Not supported
From the cloud	GaussDB primary/standby -> GaussDB distributed	Supporte d	Not supported	Supporte d
From the cloud	GaussDB primary/standby -> GaussDB primary/standby	Supporte d	Not supported	Supporte d
From the cloud	GaussDB distributed -> MySQL	Supporte d	Not supported	Not supported
From the cloud	GaussDB distributed -> Oracle	Supporte d	Not supported	Supporte d
From the cloud	GaussDB distributed -> GaussDB(DWS)	Supporte d	Not supported	Not supported
From the cloud	GaussDB distributed -> Kafka	Not supported	Not supported	Supporte d
From the cloud	GaussDB distributed -> GaussDB distributed	Supporte d	Not supported	Supporte d
From the cloud	GaussDB distributed->GaussDB primary/standby	Supporte d	Not supported	Supporte d

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Additional Column	Column Processin g
From the cloud	GaussDB(for MySQL)->MySQL	Supporte d	Supporte d	Not supported
From the cloud	GaussDB(for MySQL)- >GaussDB(DWS)	Not supported	Supporte d	Not supported
From the cloud	GaussDB(for MySQL)->CSS/ES	Supporte d	Not supported	Supporte d
From the cloud	MariaDB->MariaDB	Supporte d	Not supported	Not supported
Self- built -> Self- built	MySQL->Kafka	Not supported	Not supported	Supporte d
Self- built -> Self- built	MySQL->CSS/ES	Supporte d	Not supported	Supporte d
Self- built -> Self- built	MySQL -> GaussDB Distributed	Supporte d	Supporte d	Supporte d
Self- built -> Self- built	MySQL -> GaussDB Primary/ Standby	Supporte d	Supporte d	Supporte d
Self- built -> Self- built	Oracle -> GaussDB primary/ standby	Supporte d	Not supported	Supporte d

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Additional Column	Column Processin g
Self-built -> Self-built	Oracle -> GaussDB distributed	Supporte d	Not supported	Supporte d
Self-built -> Self-built	GaussDB primary/standby -> Kafka	Not supported	Not supported	Supporte d
Self-built -> Self-built	GaussDB distributed -> Kafka	Not supported	Not supported	Supporte d
Self-built -> Self-built	DB2 for LUW -> GaussDB primary/ standby	Supporte d	Not supported	Not supported
Self-built -> Self-built	DB2 for LUW -> GaussDB distributed	Supporte d	Not supported	Not supported

Adding Additional Columns

Step 1 On the **Process Data** page of the real-time synchronization task, click **Additional Columns**, locate the table to be processed, and click **Add** in the **Operation** column.

Figure 7-36 Additional columns

Delongest Database/Table	New name	Column Name	Operation Type	Type	Operation
gbest_order_payment_1	gbest_order_payment_1	--	--	--	Add Delete
gbest_order_payment_10	gbest_order_payment_10	--	--	--	Add Delete
gbest_order_payment_11	gbest_order_payment_11	--	--	--	Add

Step 2 In the displayed **Add** dialog box, specify the column name, operation type, and field type. Click **OK**.

Figure 7-37 Operation types

The screenshot shows a dialog box titled "Add" with a close button (X) in the top right corner. It contains two main input fields:

- * Column Name**: A text input field.
- * Operation Type**: A dropdown menu with an upward arrow on the right. The dropdown is open, showing three options:
 - Default**: The create_time column is used. The update_time column is used.
 - Expression**: The serverName@database@table column is used.
 - Value**: (partially visible, description not fully shown).

An "OK" button is located on the right side of the dialog box.

NOTE

- In many-to-one mapping scenarios, additional columns for data processing are required to avoid data conflicts.
- The following operation types are supported:
 - **Default**: Use the default value to fill in the new column.
 - Use the create_time column and update_time column as an example to fill the new column with the data creation time and data update time.
 - **Expression**: Use the **concat(_current_database, '@', _current_table)** expression to fill in the new column. You cannot manually enter an expression.
 - If you fill in the new column in **serverName@database@table** format, you need to enter a server name and then the database name and table name will be automatically filled in.
 - **Value**: Select a value, for example, synchronization time.
- You can apply the additional column information of the first editable table to all editable tables in batches.
- During MySQL to GaussDB(for MySQL) synchronization, if the number of columns in a single table exceeds 500, the number of additional columns added to the table may exceed the upper limit. As a result, the task fails.
- If **serverName@database@table** is used to add an additional column, this additional column will be used on the destination database as an implicit filtering condition for row comparison and value comparison by default.
- For a table with additional columns, the DDL operations of dropping a table and then creating a table are not supported in the incremental synchronization phase.
- During task editing in the many-to-one mapping scenario, if the new table has been synchronized, many-to-one mapping has been performed, and additional columns have been set, you need to reset the additional columns for the table. Otherwise, the additional column settings in the last synchronization are retained by default.

Step 3 Click **Next**.

----End

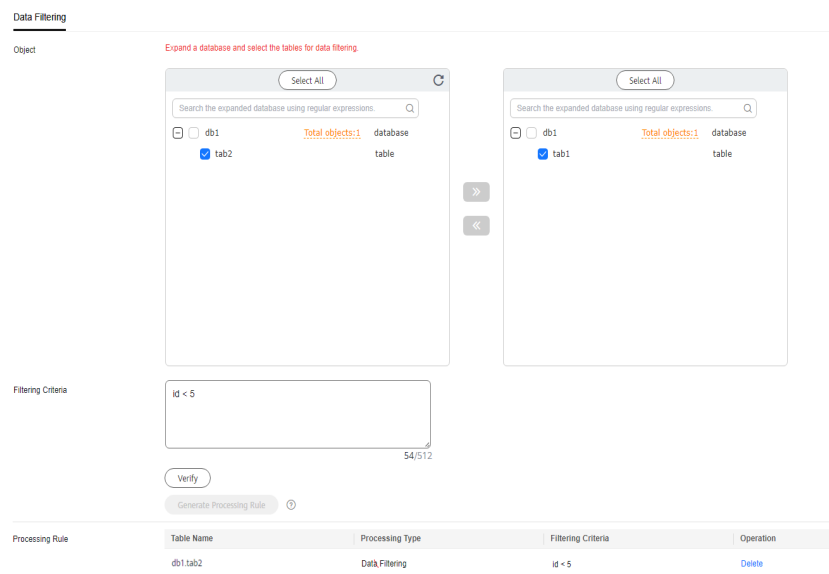
Filtering Data

After a data filtering rule is added, update the source database to ensure data consistency. For example:

- The filter criteria are met after the update. You need to continue the synchronization and perform the same update operation on the destination database. If no data is matched, the operation will be ignored, causing data inconsistency.
- The filter criteria are not met after the update. You need to continue the synchronization and perform the same update operation on the destination database.

Step 1 On the **Processing Data** page, set **Processing Type** to **Data filtering**.

Figure 7-38 Filtering data



Step 2 In the **Object** area, select the table to be processed.

Step 3 In the **Filtering Criteria** area, enter the filter criteria (only the part after WHERE in the SQL statement, for example, id=1), and click **Verify**.

 NOTE

- Each table has only one verification rule.
- Up to 512 tables can be filtered at a time. If there are more than 512 tables, perform rule verifications in batches.
- The filter expression cannot use the package, function, variable, or constant of a specific DB engine. It must comply with the general SQL standard. Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, `sid > 3` and `sname like "G %"`. A maximum of 512 characters are allowed.
- In SQL statements for setting filter criteria, keywords must be enclosed in backquotes, and the value of **datetime** (including date and time) and character string type must be enclosed in single quotation marks, for example, ``update` > '2022-07-13 00:00:00'` and `age >10, `update` ='abc'`.
- If the **TIMESTAMP** type is used as a filtering condition, the time of the character type must be set to the time value in the UTC time zone. For example, in MySQL, the **TIMESTAMP** data is stored based on the UTC time zone. You need to use the time value in the UTC time zone for comparison.
- Implicit conversion rules are not supported. Enter filtering criteria of a valid data type. For example, if column `c` of an Oracle database uses characters of the `varchar2` type, the filtering criteria must be set to `c > '10'` instead of `c > 10`.
- Filter criteria cannot be configured for large objects, such as `CLOB`, `BLOB`, and `BYTEA`.
- Filtering rules cannot be set for objects whose database names and table names contain newline characters.
- The syntax of row-level locks, such as `for update`, cannot be used as filtering criteria.
- Function operations cannot be performed on column names. If function operations are performed, data may be inconsistent.
- You are not advised to set filter criteria for fields of approximate numeric types, such as `FLOAT`, `DECIMAL`, and `DOUBLE`.
- Do not use fields containing special characters as a filter condition.
- You are advised not to perform DDL operations on columns involved in filter criteria. Otherwise, task exceptions may occur.
- You are not advised to use non-idempotent expressions or functions as data processing conditions, such as `SYSTIMESTAMP` and `SYSDATE`, because the returned result may be different each time the function is called.
- The filtering rules for a synchronized table cannot be modified.
- During data filtering for real-time synchronization with Oracle serving as the source database, the fixed-length character types `NCHAR` and `CHAR` must be matched using complete fixed-length characters.

Step 4 After the verification is successful, click **Generate Processing Rule**. The rule is displayed.

Step 5 Click **Next**.

----End

Advanced Settings for Data Filtering

If you need to query an association table, you can use the advanced settings of data processing.

Step 1 On the **Process Data** page of the real-time synchronization task, set **Processing Type** to **Data filtering**.

Step 2 In the **Object** area, select the table to be processed.

Step 3 In the **Filtering Criteria** area, specify the filtering criteria, for example, id1 in (select id from db1.tab1 where id >=3 and id <10), and click **Verify**.

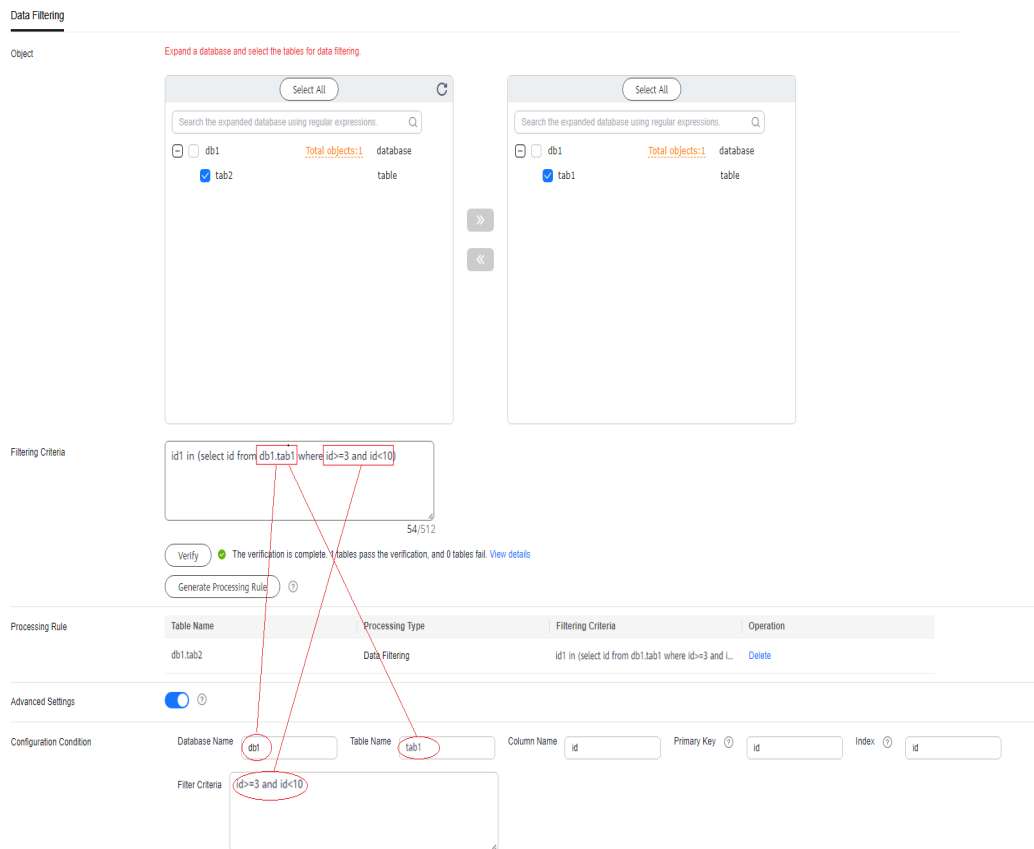
 **NOTE**

- Each table has only one verification rule.
- Up to 512 tables can be filtered at a time. If there are more than 512 tables, perform rule verifications in batches.
- The filter expression cannot use the package, function, variable, or constant of a specific DB engine. It must comply with the general SQL standard. Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, sid > 3 and sname like "G %". A maximum of 512 characters are allowed.
- Implicit conversion rules are not supported. Enter filtering criteria of a valid data type. For example, if column c of an Oracle database uses characters of the varchar2 type, the filtering criteria must be set to c > '10' instead of c > 10.
- Filter criteria cannot be configured for large objects, such as CLOB, BLOB, and BYTEA.
- Filtering rules cannot be set for objects whose database names and table names contain newline characters.
- The syntax of row-level locks, such as for update, cannot be used as filtering criteria.
- Data changes in a referenced table are not supported, which may cause data inconsistency during synchronization.
- You are not advised to set filter criteria for fields of approximate numeric types, such as FLOAT, DECIMAL, and DOUBLE.
- Do not use fields containing special characters as a filter condition.
- You are not advised to use non-idempotent expressions or functions as data processing conditions, such as SYSTIMESTAMP and SYSDATE, because the returned result may be different each time the function is called.
- During data filtering for real-time synchronization with Oracle serving as the source database, the fixed-length character types NCHAR and CHAR must be matched using complete fixed-length characters.

Step 4 After the verification is successful, click **Generate Processing Rule**. The rule is displayed.

Step 5 In the **Advanced Settings** area, specify the configuration condition and rule for the association table to help you filter data.

Figure 7-39 Advanced settings



1. In the **Configuration Condition** area, enter the association table information entered in **Step 3**.
Database Name, Table Name, Column Name, Primary Key, Index, and Filter Criteria are mandatory. If the table does not have an index, enter its primary key.
Filter Criteria is the filter condition of the association table information entered in **Step 3**.
2. Then, click **Verify**.
3. After the verification is successful, click **Generate Configuration Rule**. The rule is displayed in the **Configuration Rule** area.
To filter data in multiple association tables, repeat **Step 5**.

NOTE

Configuration rules can be deleted.

Step 6 Click **Next**.

----End

Processing Columns

Step 1 On the **Process Data** page of the real-time synchronization task, select **Processing Columns**.

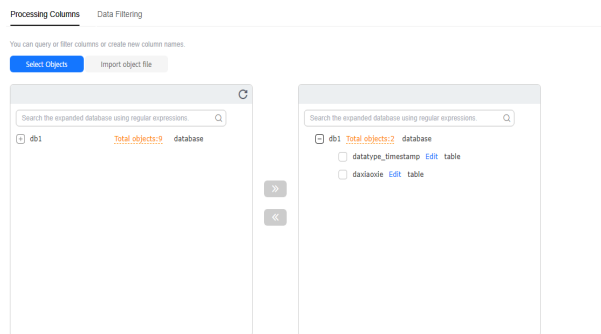
Step 2 Select a column processing mode.

 **NOTE**

Only MySQL-to-GaussDB and Oracle-to-GaussDB synchronization tasks support column processing by importing files. For other tasks, column processing is performed by selecting objects by default.

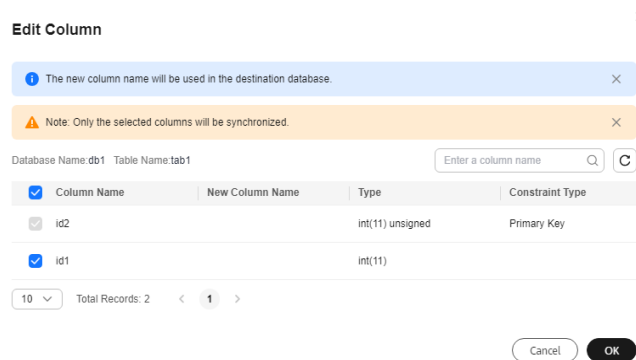
- **Select Objects**
 - a. In the **Object** area, select the objects to be processed.

Figure 7-40 Processing columns



- b. Click **Edit** to the right of the selected object.
- c. In the **Edit Column** dialog box, select the columns to be mapped and enter new column names.


Figure 7-41 Editing a column



 **NOTE**

- You can query or filter columns or create new column names.
 - After the column name is edited, the column name of the destination database is changed to the new name.
 - The new column name cannot be the same as the original column name or an existing column name.
 - Columns whose database names or table names contain newline characters cannot be mapped.
 - The column name in the synchronized table cannot be modified.
 - Only selected columns are synchronized. Newly-added columns are not included in column processing.
 - The partitioned table does not support column mapping or column filtering.
 - In the incremental phase, DDL operations cannot be performed on filtered, mapped, additional columns in a table.
 - For a table on which column filtering, column mapping, and additional column adding are performed, the DDL operations of dropping a table and then creating a table are not supported in the incremental synchronization phase.
 - If the source database is MySQL or GaussDB(for MySQL), column filtering and mapping are not supported for columns that have function-based indexes.
- d. Click **OK**.
- **Import object file**
 - a. On the **Process Data** page of the real-time synchronization task, choose **Processing Columns > Import object file**.
 - b. Click **Download Template**.

Figure 7-42 Processing columns




- c. In the downloaded Excel file, enter information about the objects to be imported.
- d. Click **Select File**. In the displayed dialog box, select the edited template.
- e. Click **Upload**.

Step 3 Click **Next**.

----End

Viewing Data Filtering Results

Step 1 On the **Data Synchronization Management** page, click the task to be processed.

Step 2 Click the **Process Data** tab to view data filtering records. Click  in the upper right corner to refresh the record list.

----End

View Column Processing

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.


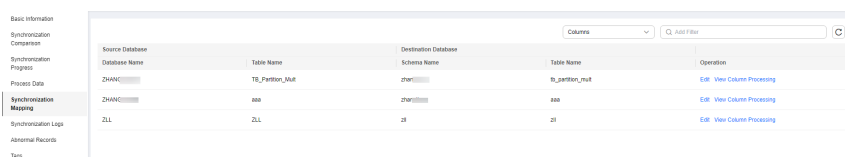
Step 2 In the navigation pane on the left, choose **Synchronization Mapping**. In the upper right corner, and select **Columns** to view column mapping records. Click  in the upper right corner to refresh the record list.

Figure 7-43 Viewing column mappings



Source Database		Destination Database		
Database Name	Table Name	Schema Name	Table Name	Operation
ZHANG	TE_Partition_Mult	zhan	te_partition_mult	Edit View Column Processing
ZHANG	aaa	zhan	aaa	Edit View Column Processing
ZLL	ZLL	ZL	ZL	Edit View Column Processing

----End

7.6 Managing Parameters

7.6.1 Changing Task Parameters

DRS allows you to change task parameters based on service requirements.

Constraints

- You can change parameters of a synchronization task in the **Configuration**, **Full**, **Full synchronization failed**, **Incremental**, **Incremental synchronization failed**, or **Paused** state.
- After some parameters of a task are changed, the changes take effect only after you restart the task by referring to [Restarting a Synchronization Task](#). In addition, operations for resetting a task, resuming a task, rebuilding a task, performing primary/standby switchover, or changing the flow control in the full synchronization phase will restart task processes for the parameter changes to take effect.

- You can change parameters only for the synchronization tasks of the following engines: For details about common task parameters, see [Task Parameter Description](#).
 - MySQL->MySQL
 - MySQL->PostgreSQL
 - MySQL -> GaussDB Distributed
 - MySQL -> GaussDB Primary/Standby
 - MySQL->GaussDB(DWS)
 - MySQL->GaussDB(for MySQL)
 - MySQL->Kafka
 - MySQL->CSS/ES
 - MySQL->Oracle
 - Oracle->MySQL
 - Oracle->PostgreSQL
 - Oracle -> GaussDB Distributed
 - Oracle -> GaussDB Primary/Standby
 - Oracle->GaussDB(for MySQL)
 - Oracle->DDM
 - Oracle->Kafka
 - Microsoft SQL Server->MySQL
 - Microsoft SQL Server->PostgreSQL
 - Microsoft SQL Server -> GaussDB Distributed
 - Microsoft SQL Server -> GaussDB Primary/Standby
 - Microsoft SQL Server->GaussDB(for MySQL)
 - Microsoft SQL Server->Kafka
 - Microsoft SQL Server->Microsoft SQL Server
 - DB2 for LUW -> GaussDB Distributed
 - DB2 for LUW -> GaussDB Primary/Standby

Procedure

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task.
- Step 2** Click the task name. The **Basic Information** page is displayed.
- Step 3** Choose **Parameters** on the left pane. On the **Parameters** tab page, change the parameters of the current synchronization task.

Figure 7-44 Parameters

Parameter Name	Effective upon Restart	Value	Allowed Values	Description
apply_thread_num	Yes	8	1-16	Number of threads for writing data to the destination da...
relay_socket_timeout	Yes	30000	30000-2147483647	Timeout interval for obtaining logs from the source data...
shard_bin_num	Yes	4	1-8	Number of sharding threads in the source database du...
datanode_source_socket_timeout	Yes	120000	120000-2147483647	Timeout interval for obtaining data from the source dat...
read_bin_num	Yes	4	1-8	Number of threads for reading data from the source da...
increment_writer_num	Yes	64	1-128	Number of threads for writing data to the destination da...

Step 4 Click **Save** or **Cancel**.

- To save your changes, click **Save**. In the displayed dialog box, click **Yes**.
- To cancel your changes, click **Cancel**. In the displayed dialog box, click **Yes**.

Step 5 Restart the task for the changes to take effect. You can click **Restart** in the upper right corner or restart the task by referring to [Restarting a Synchronization Task](#).

In addition, operations for resetting a task, resuming a task, rebuilding a task, performing primary/standby switchover, or changing the flow control in the full synchronization phase will restart task processes for the parameter changes to take effect.

----End

Task Parameter Description

Table 7-12 Task parameters

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
applier_thread_num	Number of threads for writing data to the destination database during the full synchronization phase.	The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-4 • Small: 1-4 • Medium: 1-8 • Large: 1-18 • Ultra-large: 1-24 	The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 2 • Small: 2 • Medium: 4 • Large: 8 • Ultra-large: 12 	int	Yes	<ul style="list-style-type: none"> • MySQL->MySQL • MySQL->PostgreSQL • MySQL->GaussDB Distributed • MySQL->GaussDB Primary / Standby • MySQL->GaussDB(DWS) • MySQL->GaussDB(for MySQL) • MySQL->Kafka • MySQL->CSS/ES • MySQL->Oracle

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						<ul style="list-style-type: none"> • Microsoft SQL Server->MySQL • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->GaussDB(for MySQL) • Microsoft SQL Server->Micros

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						oft SQL Server <ul style="list-style-type: none"> • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
read_task_num	Number of threads for reading data from the source database during the full synchronization phase.	The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-2 • Small: 1-2 • Medium: 1-4 • Large: 1-8 • Ultra-large: 1-12 	The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1 • Small: 1 • Medium: 2 • Large: 4 • Ultra-large: 6 	int	Yes	<ul style="list-style-type: none"> • MySQL->MySQL • MySQL->PostgreSQL • MySQL->GaussDB Distributed • MySQL->GaussDB Primary / Standby • MySQL->GaussDB(DWS) • MySQL->GaussDB(for MySQL) • MySQL->Kafka • MySQL->CSS/ES • MySQL->Oracle • Microsoft SQL Server->MySQL

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						<ul style="list-style-type: none"> • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->GaussDB(for MySQL) • Microsoft SQL Server->Microsoft SQL Server • DB2 for LUW -> GaussDB

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						B Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
shard_task_num	Number of sharding threads in the source database during the full synchronization phase.	The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-2 • Small: 1-2 • Medium: 1-4 • Large: 1-8 • Ultra-large: 1-8 	The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1 • Small: 1 • Medium: 2 • Large: 4 • Ultra-large: 4 	int	Yes	<ul style="list-style-type: none"> • MySQL->MySQL • MySQL->PostgreSQL • MySQL->GaussDB Distributed • MySQL->GaussDB Primary / Standby • MySQL->GaussDB(DWS) • MySQL->GaussDB(for MySQL) • MySQL->Kafka • MySQL->CSS/ES • MySQL->Oracle • Oracle->MySQL

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						<ul style="list-style-type: none"> ● Oracle->PostgreSQL ● Oracle -> GaussDB Distributed ● Oracle -> GaussDB Primary / Standby ● Oracle->GaussDB(for MySQL) ● Oracle->DDM ● Microsoft SQL Server->MySQL ● Microsoft SQL Server->PostgreSQL ● Microsoft SQL Server -> GaussD

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						B Distributed <ul style="list-style-type: none"> • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->Gauss DB(for MySQL) • Microsoft SQL Server->Microsoft SQL Server • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
increment_writer_num	Number of threads for writing data to the destination database during the incremental synchronization phase.	The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-8 • Small: 1-16 • Medium: 1-32 • Large: 1-128 • Ultra-large: 1-128 	The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 4 • Small: 8 • Medium: 16 • Large: 64 • Ultra-large: 64 	int	Yes	<ul style="list-style-type: none"> • MySQL->MySQL • MySQL->PostgreSQL • MySQL->GaussDB Distributed • MySQL->GaussDB Primary / Standby • MySQL->GaussDB(DWS) • MySQL->GaussDB(for MySQL) • MySQL->Kafka • MySQL->CSS/ES • MySQL->Oracle

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
relay_socket_timeout	Timeout interval for obtaining logs from the source database during the incremental synchronization phase.	30000-2147483647 ms	30000 ms	int	Yes	<ul style="list-style-type: none"> • MySQL->MySQL • MySQL->PostgreSQL • MySQL->GaussDB Distributed • MySQL->GaussDB Primary / Standby • MySQL->GaussDB(DWS) • MySQL->GaussDB(for MySQL) • MySQL->Kafka • MySQL->CSS/ES • MySQL->Oracle

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
datamove_source_socket_timeout	Timeout interval for obtaining data from the source database during the full synchronization phase.	<p>The value range depends on the DRS task engine.</p> <ul style="list-style-type: none"> MySQL -> MySQL: 120000-214748 3647 ms MySQL -> PostgreSQL: 120000-214748 3647 ms MySQL -> GaussDB Distributed: 120000-214748 3647 ms MySQL -> GaussDB Primary/Standby: 120000-214748 3647 ms MySQL -> GaussDB(DWS): 120000-214748 3647 ms MySQL -> GaussDB (for MySQL): 120000-214748 3647 ms MySQL -> Kafka: 120000-214748 3647 ms 	<p>The value range depends on the DRS task engine.</p> <ul style="list-style-type: none"> MySQL -> MySQL: 120000 ms MySQL -> PostgreSQL: 120000 ms MySQL -> GaussDB Distributed: 120000 ms MySQL -> GaussDB Primary/Standby: 120000 ms MySQL -> GaussDB(DWS): 120000 ms 	int	Yes	<ul style="list-style-type: none"> MySQL->MySQL MySQL->PostgreSQL MySQL->GaussDB Distributed MySQL->GaussDB Primary/Standby MySQL->GaussDB(DWS) MySQL->GaussDB(for MySQL) MySQL->Kafka MySQL->CSS/ES MySQL->Oracle Oracle->MySQL

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
		<ul style="list-style-type: none"> MySQL -> CSS/ES: 120000-214748 3647 ms MySQL -> Oracle: 120000-214748 3647 ms Oracle -> MySQL: 600000-214748 3647 ms Oracle -> PostgreSQL: 600000-214748 3647 ms Oracle -> GaussDB Distributed: 600000-214748 3647 ms Oracle -> GaussDB Primary/ Standby: 600000-214748 3647 ms Oracle -> GaussDB(for MySQL): 600000-214748 3647 ms Oracle -> DDM: 600000-214748 3647 ms 	<ul style="list-style-type: none"> MySQL -> GaussDB (for MySQL): 120000 ms MySQL -> Kafka: 120000 ms MySQL -> CSS/ES: 120000 ms MySQL->Oracle : 120000 ms Oracle -> MySQL: 600000 ms Oracle -> PostgreSQL: 600000 ms Oracle -> GaussDB Distributed: 600000 ms Oracle -> GaussDB 			<ul style="list-style-type: none"> Oracle->PostgreSQL Oracle -> GaussDB Distributed Oracle -> GaussDB Primary / Standby Oracle->Gauss DB(for MySQL) Oracle->DDM Oracle->Kafka Microsoft SQL Server->MySQL Microsoft SQL Server->PostgreSQL Microsoft SQL Server -

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
		<ul style="list-style-type: none"> • Oracle -> Kafka: 600000-214748 3647 ms • Microsoft SQL Server -> MySQL: 360000-214748 3647 ms • Microsoft SQL Server -> PostgreSQL: 360000-214748 3647 ms • Microsoft SQL Server -> GaussDB Distributed: 360000-214748 3647 ms • Microsoft SQL Server -> GaussDB Primary/ Standby: 360000-214748 3647 ms • Microsoft SQL Server -> GaussDB(for MySQL): 360000-214748 3647 ms • Microsoft SQL Server -> Kafka: 	Primary/ Standby: 600000 ms <ul style="list-style-type: none"> • Oracle -> GaussDB(for MySQL): 600000 ms • Oracle -> DDM: 600000 ms • Oracle -> Kafka: 600000 ms • Microsoft SQL Server -> MySQL: 360000 ms • Microsoft SQL Server -> PostgreSQL: 360000 ms • Microsoft SQL Server -> GaussDB Distributed 			> GaussDB Distributed <ul style="list-style-type: none"> • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->GaussDB(for MySQL) • Microsoft SQL Server->Kafka • Microsoft SQL Server->Microsoft SQL Server

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
		360000-214748 3647 ms <ul style="list-style-type: none"> Microsoft SQL Server -> Microsoft SQL Server: 360000-214748 3647 ms 	ed: 360000 ms <ul style="list-style-type: none"> Microsoft SQL Server -> GaussDB Primary/ Standby: 360000 ms Microsoft SQL Server -> GaussDB(for MySQL): 360000 ms Microsoft SQL Server -> Kafka: 360000 ms Microsoft SQL Server -> Microsoft SQL Server: 360000 ms 			

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
shard_length	Rows per shard during full synchronization	520000-100000000 or 0	520000	int	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM • Microsoft SQL Server->MySQL • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						<ul style="list-style-type: none"> > GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->GaussDB(for MySQL) • Microsoft SQL Server->Microsoft SQL Server • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						/ Standby
datamove_fetchsize	The size of the data reads from the source database during a full synchronization.	10-10000	The value range depends on the DRS task engine. <ul style="list-style-type: none"> • Oracle -> MySQL: 1000 • Oracle->PostgreSQL : 1000 • Oracle -> GaussDB Distributed: 10000 • Oracle -> GaussDB Primary/ Standby: 10000 • Oracle->GaussDB(for MySQL) : 1000 • Oracle->DDM : 1000 • Oracle->Kafka : 1000 	int	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
datamove_copy_mode	<p>The write mode used when writing data to the destination database during an incremental synchronization. The write performance of COPY is higher than that of INSERT.</p> <p>true indicates the COPY mode, and false indicates the INSERT mode.</p>	true/false	true	boolean	Yes	<ul style="list-style-type: none"> • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						/ Standby • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
datamove_split_partition	Whether to shard a partition table during full synchronization. true: Partition tables in the source database are sharded for parallel synchronization to improve the synchronization speed of large partitions. false: Partition tables are not sharded.	true/false	true	boolean	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
read_log_num	Number of threads for obtaining logs from the source database during incremental synchronization.	1-16	2	int	Yes	<ul style="list-style-type: none"> ● Oracle->MySQL ● Oracle->PostgreSQL ● Oracle -> GaussDB Distributed ● Oracle -> GaussDB Primary / Standby ● Oracle->GaussDB(for MySQL) ● Oracle->DDM ● Oracle->Kafka

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
logminer_fetchsize	The size of the data reads when using LogMiner in scenarios where the source database is Oracle and incremental log read mode is set to LogMiner.	10-10000	1000	int	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM • Oracle->Kafka

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
deal_hot_table	Whether to accelerate the replay of hot tables with frequent DML operations in the source database during incremental synchronization. true: The replay of hot tables is accelerated; false: The replay of hot tables is not accelerated.	true/false	false	boolean	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM • Microsoft SQL Server->MySQL • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						<ul style="list-style-type: none"> > GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary / Standby • Microsoft SQL Server->GaussDB(for MySQL) • Microsoft SQL Server->Microsoft SQL Server • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						/ Standby
skip_all_ddl	Whether to ignore all DDL statements during incremental synchronization. If there are a lot of irrelevant DDL statements, enable this function to improve the incremental synchronization performance. true: All DDL statements are ignored. false: All DDL statements are not ignored.	true/false	false	boolean	Yes	<ul style="list-style-type: none"> • Oracle->MySQL • Oracle->PostgreSQL • Oracle->GaussDB Distributed • Oracle->GaussDB Primary / Standby • Oracle->GaussDB(for MySQL) • Oracle->DDM • Oracle->Kafka

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
replace_invalid_time	Whether to replace the source time type with the default time when the source time type is invalid after being synchronized to the destination database during full+incremental synchronization. Default timestamp: 1970-01-01 00:00:00; default date: 1970-01-01; default time: 00:00:00. true: the source	true/false	false	boolean	Yes	<ul style="list-style-type: none"> • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
	time type is replaced with the default time; false : the source time type is not replaced with the default time.					/ Standby • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
replace_0x00	Whether to replace 0x00 characters that are not supported by the destination database with spaces during full +incremental synchronization. true: 0x00 characters are replaced with spaces; false: 0x00 characters are not replaced with spaces.	true/false	true	boolean	Yes	<ul style="list-style-type: none"> • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
						/ Standby • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
replace_0xefbfb	Whether to replace abnormal characters with the default value '?' during full +incremental synchronization. DRS sets the character set of the JDBC client to UTF-8. During the query in the source database, abnormal bytes are converted to the UTF-8-encoded exception code EFBFBD. true: Abnormal code in UTF-8	true/false	false	boolean	Yes	<ul style="list-style-type: none"> • Oracle->PostgreSQL • Oracle -> GaussDB Distributed • Oracle -> GaussDB Primary / Standby • Microsoft SQL Server->PostgreSQL • Microsoft SQL Server -> GaussDB Distributed • Microsoft SQL Server -> GaussDB Primary

Parameter	Description	Value Range	Default Value	Type	Restart Required	Applicable Data Flow
	encoding is replaced with '?'. false: Abnormal code is not replaced.					/ Standby • DB2 for LUW -> GaussDB Distributed • DB2 for LUW -> GaussDB Primary / Standby

7.6.2 Viewing Parameter Change History

DRS allows you to view the change history of task parameters to meet service requirements.

Procedure

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task.
- Step 2** Click the task name. The **Basic Information** page is displayed.
- Step 3** Choose **Parameters** on the left pane. On the **Change History** tab page, view the change history of parameters of the current synchronization task.

By default, the parameter change history of the last seven days can be queried.

Figure 7-45 Change History

Parameter Name	Original Value	New Value	Change Status	Change Time	Apply or Not	Application Time
share_task_num	4	6	Successful	Apr 24, 2024 11:34:11 GMT+08:00	No	--

----End

7.7 Task Life Cycle

7.7.1 Viewing Task Details

View the information about the synchronization tasks and synchronization instances. This section describes how to view details about a synchronization task you have created.

Prerequisites

You have logged in to the DRS console.

Procedure

NOTE

In the task list, only tasks created by the current login user are displayed. Tasks created by different users of the same tenant are not displayed.

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 On the displayed **Basic Information** page, view details about the migration task.

You can view information about the task, synchronization instance, and synchronization.

----End

7.7.2 Modifying Task Information

After a synchronization task is created, you can modify task information to identify different tasks.

The following task information can be edited:

- Task name
- Description
- SMN topic
- Synchronization delay threshold
- Number of days when an abnormal task is stopped

- Task start time

Prerequisites

You have logged in to the DRS console.

Procedure




- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, locate the information to be modified in the **Task Information** area.
- You can click  to modify the task name, SMN topic, delay threshold, the time to stop abnormal tasks, and description.
 - To submit the change, click .
 - To cancel the change, click .

Table 7-13 Task information

Task Information	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain the following special characters: !<>&'\"
SMN Topic	You can apply for a topic on the SMN console and add a subscription. For details, see Simple Message Notification User Guide .
Synchronization Delay Threshold	The delay ranges from 0s to 3600s. NOTE If the delay threshold is set to 0, no notifications will be sent to the recipient.
Stop Abnormal Tasks After	The value must range from 14 to 100. The default value is 14. NOTE You can set this parameter only for pay-per-use tasks.

- You can modify the task start time only when the task is in the **Pending start** status.

In the **Task Information** area, click **Modify** in the **Scheduled Start Time** field. On the displayed page, specify the scheduled start time and click **OK**.


- Step 3** View the change result on the **Basic Information** tab.

----End

Configuring Exception Notifications

- Step 1** On the **Data Synchronization Management** page, select the task to be configured.
- Step 2** Click **Batch Operations** in the upper left corner and choose **Configure Exception Notification**.

Figure 7-46 Batch Operations



- Step 3** In the displayed dialog box, enter the configuration information and click **Yes** to submit the configuration task.

----End

7.7.3 Modifying Connection Information

During the synchronization, you may change the password of the source or destination database. As a result, the data synchronization, data comparison, task resuming, resetting, object editing, and stopping may fail. In this case, you need to change the password on the DRS console and resume the task.

You can modify the following synchronization information:

- Database password
- Database IP address
- Database port
- Database username

Constraints

- The database connection password can be changed for all DRS tasks.
- You can change the IP address, port, and username during the incremental synchronization phase only for a DRS task with MySQL or GaussDB(for MySQL) serving as the source and IP address entered for the connection test. If the IP address, port number, or username changes due to some operations on the source database, you can use this function to update the information.
- The function of changing an IP address applies to the scenario where the IP address of the source database changes. The IP addresses before and after the change must belong to the same data instance. Otherwise, the task may fail or data may be inconsistent.
- After the connection information is changed, the change takes effect immediately, and the data in the destination database is not cleared.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click **Modify Connection Details** in the **Connection Information** area.
- Step 3** In the displayed dialog box, change the database information such as passwords of the source and destination databases and click **OK**.
- Step 4** If a task is in the **Failed** state, return to the task list page after the change is complete and click **Resume** in the **Operation** column to continue the synchronization task.

----End

7.7.4 Modifying the Flow Control Mode

You can choose whether to control the flow. DRS allows you to change the flow control mode after a task is created. Currently, only the real-time synchronization scenarios listed in [Real-time Synchronization Scenarios That Support Flow Control](#) support this function.

Constraints

- The flow control mode limits the maximum traffic speed in seconds. The actual statistical value may be lower than the flow rate because the statistical value may decrease due to network fluctuation.
- The flow control mode takes effect only in the full synchronization phase.
- After the traffic rate is modified in the incremental migration phase, the modification takes effect when the task enters the full migration phase again. For example, if the traffic rate is modified and a synchronization object is added to the task, the modification takes effect in the full synchronization phase of the task.

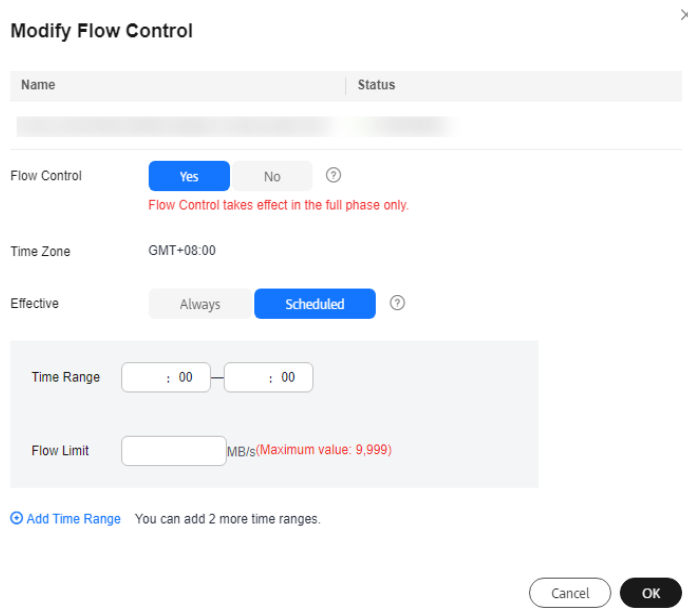
Prerequisites

- You have logged in to the DRS console.
- A synchronization task has been created.

Method 1

- Step 1** In the **Flow Control Information** area on the **Basic Information** tab, click **Modify**.
- Step 2** In the displayed dialog box, modify the settings.

Figure 7-47 Modifying the Flow Control Mode




----End

Method 2

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and choose **More > Speed** or **Speed** in the **Operation** column.

Figure 7-48 Task List



Step 2 In the displayed dialog box, modify the settings.

Figure 7-49 Modifying the Flow Control Mode

Modify Flow Control ×

Name	Status

Flow Control Yes No ?
Flow Control takes effect in the full phase only.

Time Zone GMT+08:00

Effective Always Scheduled ?

Time Range

Flow Limit MB/s (Maximum value: 9,999)

[Add Time Range](#) You can add 2 more time ranges.

----End

Real-time Synchronization Scenarios That Support Flow Control

To the cloud

- MySQL->MySQL
- MySQL->GaussDB(for MySQL)
- MySQL -> GaussDB primary/standby
- MySQL -> GaussDB Distributed
- MySQL->PostgreSQL
- MySQL->GaussDB(DWS)
- MySQL->MariaDB
- DDM->MySQL
- DDM->DDM
- DDM->GaussDB(DWS)
- Oracle->MySQL
- Oracle->GaussDB(for MySQL)
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- Oracle->GaussDB(DWS)
- Oracle->DDM
- Oracle->PostgreSQL
- PostgreSQL->PostgreSQL
- PostgreSQL->GaussDB(DWS)

- PostgreSQL -> GaussDB Primary/Standby
- PostgreSQL -> GaussDB Distributed
- TiDB->GaussDB(for MySQL)
- Microsoft SQL Server->MySQL
- Microsoft SQL Server->GaussDB(for MySQL)
- Microsoft SQL Server->GaussDB(DWS)
- Microsoft SQL Server -> GaussDB primary/standby
- Microsoft SQL Server -> GaussDB distributed
- Microsoft SQL Server->Microsoft SQL Server
- Microsoft SQL Server->PostgreSQL
- MongoDB->DDS
- MariaDB->MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- GaussDB(for MySQL)->GaussDB(for MySQL)
- Cassandra->GeminiDB Cassandra
- Dynamo->GeminiDB Dynamo

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- MySQL->MariaDB
- DDM->MySQL
- DDM->Oracle
- GaussDB(for MySQL)->MySQL
- GaussDB(for MySQL)->CSS/ES
- GaussDB(for MySQL)->Oracle
- GaussDB(for MySQL)->Kafka
- GaussDB primary/standby -> MySQL
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> GaussDB(DWS)
- GaussDB primary/standby -> GaussDB distributed
- GaussDB primary/standby -> GaussDB primary/standby
- GaussDB distributed -> MySQL
- GaussDB distributed -> Oracle
- GaussDB distributed -> GaussDB(DWS)
- GaussDB distributed -> GaussDB distributed
- GaussDB distributed->GaussDB primary/standby
- MariaDB->MariaDB

- PostgreSQL->PostgreSQL

Self-built -> Self-built

- MySQL->CSS/ES
- MySQL->Kafka
- Oracle->Kafka
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> GaussDB primary/standby
- GaussDB distributed -> Oracle
- GaussDB distributed -> GaussDB distributed




7.7.5 Changing the Synchronization Mode

DRS allows you to change the synchronization mode after a task is created.

Constraints

- You can change the synchronization mode only for MySQL-to-GaussDB(for MySQL) synchronization tasks.
- The synchronization mode can be changed from incremental to full +incremental and from full+incremental to incremental.
- You can change the synchronization mode for a task in the **Configuration**, **Incremental**, or **Incremental failed** state.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** In the **Synchronization Instance Details** area on the **Basic Information** page, click  next to **Synchronization Mode** to change the synchronization mode.
 - To submit the change, click .
 - To cancel the change, click .
- Step 3** After the synchronization task is submitted, check the task synchronization mode on the **Data Synchronization Management** page.

----End

7.7.6 Editing a Synchronization Task

For a synchronization task that has been created but not started, DRS allows you to edit the configuration information of the task, including the source and destination database details. For synchronization tasks in the following statuses, you can edit and submit the tasks again.

- Creating
- Configuration

 NOTE

For an incremental synchronization task, DRS allows you to modify synchronization objects. For details, see [Editing Synchronization Objects](#).

Prerequisites

You have logged in to the DRS console.

Method 1

- Step 1** In the task list on the **Data Synchronization Management** page, locate the target task and click **Edit** in the **Operation** column.
- Step 2** On the **Configure Source and Destination Databases** page, enter information about the source and destination databases and click **Next**.
- Step 3** On the **Set Synchronization Task** page, select synchronization objects and click **Next**.
- Step 4** On the **Check Task** page, check the synchronization task.
- Step 5** On the **Confirm Task** page, specify **Start Time**, confirm that the configured information is correct, select the check box before the agreement, and click **Next**.
- Step 6** After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

----End

Method 2

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task.
- Step 2** On the displayed page, click **edit this task** to go to the **Configure Source and Destination Databases** page.
- Step 3** Perform [Step 2](#) to [Step 6](#) from method 1.

----End

7.7.7 Resuming a Synchronization Task

A fault may occur during the synchronization due to external factors, such as insufficient storage space. After the fault is rectified based on the synchronization log information, you can resume the synchronization.

You can resume synchronization tasks in any of the following statuses:

- Synchronization failed
- Paused

 **NOTE**

- If the synchronization task fails due to non-network problems, the system will automatically resume the task three times by default. If the failure persists, you can resume the task manually.
- If the synchronization fails due to network problems, the system will automatically resume the task until the synchronization is restored.

Prerequisites

You have logged in to the DRS console.

Method 1

In the task list on the **Data Synchronization Management** page, locate the target task and click **Resume** in the **Operation** column.

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 On the displayed page, click the **Synchronization Progress** tab, and click **Resume** in the upper right corner.

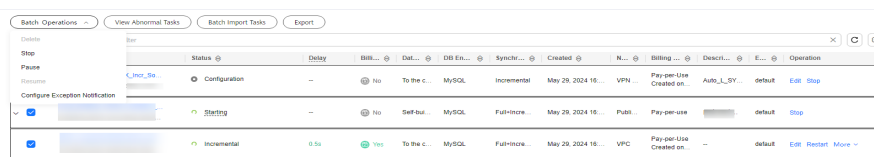
----End

Resuming Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be resumed.

Step 2 Click **Batch Operations** in the upper left corner and choose **Resume**.

Figure 7-50 Batch Operations



Stop	Phase	Status	Priority	Batch	DB EN...	Synchron...	Created	N...	Billing...	Descri...	E...	Operation
Resume	Configuration	Configuration	No	To the c.	MySQL	Incremental	May 28, 2024 16...	VPN	Pay-per-use	Auto_I_SY...	default	Edit Stop
<input checked="" type="checkbox"/>	Starting	Starting	No	Self-bui.	MySQL	Full-incre...	May 28, 2024 16...	Publi.	Pay-per-use		default	Stop
<input checked="" type="checkbox"/>	Incremental	Incremental	0.5s	Yes	To the c.	MySQL	Full-incre...	May 28, 2024 16...	VPC	Pay-per-use	Created on...	default Edit Restart More

Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

7.7.8 Skipping Data and Resuming a Synchronization Task

Data synchronization may be stopped due to certain reasons. After the problem is resolved, you can directly skip the data written to the source database during the pause and resume data transfer. Only the incremental synchronization tasks from DDM to Oracle support this function.

Prerequisites

- You have logged in to the DRS console.

- The task is paused.

Procedure

In the task list on the **Data Synchronization Management** page, locate the target task and click **Jump Resume** in the **Operation** column.

7.7.9 Pausing a Synchronization Task

DRS allows you to pause real-time synchronization tasks. For details about the synchronization scenarios where synchronization tasks can be paused, see [Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Paused](#).

Prerequisites

- You have logged in to the DRS console.

Pausing a Task

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Pause** in the **Operation** column.

Step 2 In the displayed **Pause Task** dialog box, select **Pause log capturing** and click **Yes**.

NOTE

- When a task in incremental state is paused, only the replay or capture and replay of incremental data is paused. Before database cutover, stop the task.
- After you select **Pause log capturing**, the DRS instance will no longer communicate with the source and destination databases. If the pause duration is too long, the task may fail to be resumed because the logs required by the source database expire. You are not advised pausing a task for more than 24 hours. For details, check the corresponding log configuration.
- After the task is paused, its status becomes **Paused**.
- You can use the resumable transfer function to continue the synchronization task.

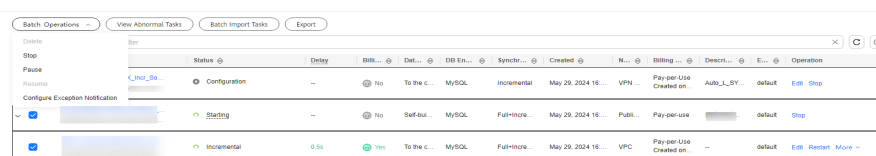
----End

Pausing Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be paused.

Step 2 Click **Batch Operations** in the upper left corner and choose **Pause**.

Figure 7-51 Batch Operations



Status	Delay	Error	Data	DB Engine	Synchronization Type	Created Time	Billing	Operation
Configuration	--	No	To the c.	MySQL	incremental	May 29, 2024 18:...	VPN Pay-per-Use Created on...	Auto_L_SY... default Edit Stop
Starting	--	No	Self-bui...	MySQL	Full-incre...	May 29, 2024 18:...	Publ... Pay-per-use	default Stop
Incremental	0.5s	Yes	To the c.	MySQL	Full-incre...	May 29, 2024 18:...	VPC Pay-per-Use Created on...	default Edit Restart More

Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Paused

The following tasks can be paused during incremental synchronization:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL -> GaussDB primary/standby
 - MySQL -> GaussDB Distributed
 - MySQL->PostgreSQL
 - MySQL->MariaDB
 - PostgreSQL->PostgreSQL
 - PostgreSQL->GaussDB(DWS)
 - PostgreSQL -> GaussDB Primary/Standby
 - PostgreSQL -> GaussDB Distributed
 - DDM->MySQL
 - DDM->DDM
 - DDM->GaussDB(DWS)
 - Oracle->GaussDB(DWS)
 - Oracle->PostgreSQL
 - Oracle->GaussDB(for MySQL)
 - Oracle -> GaussDB primary/standby
 - Oracle -> GaussDB distributed
 - Oracle->MySQL
 - Oracle->DDM
 - DB2 for LUW -> GaussDB primary/standby
 - DB2 for LUW -> GaussDB distributed
 - DB2 for LUW->GaussDB(DWS)
 - TiDB->GaussDB(for MySQL)
 - Microsoft SQL Server->MySQL
 - Microsoft SQL Server->GaussDB(for MySQL)
 - Microsoft SQL Server->GaussDB(DWS)
 - Microsoft SQL Server -> GaussDB primary/standby
 - Microsoft SQL Server -> GaussDB distributed
 - Microsoft SQL Server->Microsoft SQL Server
 - Microsoft SQL Server->PostgreSQL
 - MongoDB->DDS
 - MariaDB->MariaDB
 - MariaDB->MySQL
 - MariaDB->GaussDB(for MySQL)
 - GaussDB(for MySQL)->GaussDB(for MySQL)

- Cassandra->GeminiDB Cassandra
- Dynamo->GeminiDB Dynamo
- From the cloud
 - MySQL->MySQL
 - MySQL->CSS/ES
 - MySQL->Oracle
 - MySQL->Kafka
 - MySQL->MariaDB
 - DDM->MySQL
 - DDM->Oracle
 - DDM->Kafka
 - DDS->MongoDB
 - DDS->Kafka
 - GaussDB(for MySQL)->MySQL
 - GaussDB(for MySQL)->Kafka
 - GaussDB(for MySQL)->CSS/ES
 - GaussDB(for MySQL)->Oracle
 - GaussDB(for MySQL)->GaussDB(DWS)
 - GaussDB primary/standby -> MySQL
 - GaussDB primary/standby -> Oracle
 - GaussDB primary/standby -> Kafka
 - GaussDB distributed -> MySQL
 - GaussDB distributed -> Oracle
 - GaussDB distributed -> Kafka
 - PostgreSQL->PostgreSQL
 - PostgreSQL->Kafka
 - MariaDB->MariaDB
 - Microsoft SQL Server->Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - MySQL->CSS/ES
 - MySQL -> GaussDB Primary/Standby
 - MySQL -> GaussDB Distributed
 - Oracle-> Kafka
 - Oracle -> GaussDB primary/standby
 - Oracle -> GaussDB distributed
 - GaussDB primary/standby -> Oracle
 - GaussDB primary/standby -> Kafka
 - GaussDB distributed -> Oracle
 - GaussDB distributed -> Kafka

- PostgreSQL->Kafka
- DB2 for LUW -> GaussDB primary/standby
- DB2 for LUW -> GaussDB distributed
- Microsoft SQL Server->Kafka

In addition, the following tasks can be paused during full synchronization:

- MySQL->MySQL
- MySQL->GaussDB(for MySQL)
- MySQL->GaussDB(DWS)
- MySQL->CSS/ES
- MySQL->Kafka
- GaussDB(for MySQL)->MySQL
- GaussDB(for MySQL)->GaussDB(for MySQL)
- GaussDB(for MySQL)->Kafka
- GaussDB(for MySQL)->GaussDB(DWS)
- GaussDB(for MySQL)->CSS/ES
- Oracle->MySQL
- Oracle->PostgreSQL
- Oracle->GaussDB(for MySQL)
- Oracle->DDM
- Oracle->GaussDB(DWS)
- TiDB->GaussDB(for MySQL)
- DDM->DDM
- DDM->MySQL
- DDM->Oracle
- DDM->GaussDB(DWS)
- PostgreSQL->PostgreSQL
- PostgreSQL->GaussDB(DWS)
- MongoDB->DDS
- DDS->MongoDB
- Cassandra->GeminiDB Cassandra
- GaussDB Primary/Standby -> MySQL
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> GaussDB(DWS)
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Primary/Standby -> GaussDB Distributed
- GaussDB Distributed -> MySQL
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> GaussDB(DWS)
- GaussDB Distributed -> GaussDB Distributed
- GaussDB Distributed -> GaussDB Primary/Standby

7.7.10 Resetting a Synchronization Task

During real-time synchronization, you can reset the synchronization tasks in one of the following statuses so that you do not need to configure the tasks again.

- Paused
- Failed

For details about the synchronization scenarios where synchronization tasks can be reset, see [Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Reset](#).

NOTE

Resetting a task does not clear the destination database. You can determine whether to clear the destination database based on your service requirements.

- **Full** and **full+incremental** tasks: To ensure data consistency before and after synchronization, manually clear the destination database and reset the task. After the task is reset, full synchronization is performed again. You do not need to configure the task again.
- **Incremental** tasks: Only incremental data is synchronized. You can directly reset the task without clearing the destination database.
- If the synchronization objects contain tables without primary keys, temporary tables prefixed with **drs_** may exist after resetting a task during full synchronization. You need to manually delete the temporary tables.

Prerequisites

You have logged in to the DRS console.

Method 1

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Reset** in the **Operation** column.

Step 2 In the displayed dialog box, check the synchronization task again.

NOTE

If a many-to-one synchronization task fails to be reset, click the name of the failed subtask in the failure details to view the failure cause of the task.

Step 3 After the check is complete and the check success rate is 100%, click **Start** to submit the synchronization task again.

----End

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 On the displayed page, click the **Synchronization Progress** tab, and click **Reset** in the upper right corner.

Step 3 Perform [Step 2](#) to [Step 3](#) from method 1.

----End

Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Reset

To the cloud

- MySQL->MySQL
- MySQL->PostgreSQL
- MySQL->GaussDB(for MySQL)
- MySQL -> GaussDB primary/standby
- MySQL -> GaussDB Distributed
- MySQL->GaussDB(DWS)
- MySQL->MariaDB
- PostgreSQL->PostgreSQL
- PostgreSQL->GaussDB(DWS)
- PostgreSQL -> GaussDB Primary/Standby
- PostgreSQL -> GaussDB Distributed
- DDM->MySQL
- DDM->DDM
- DDM->GaussDB(DWS)
- Oracle->MySQL
- Oracle->GaussDB(for MySQL)
- Oracle->GaussDB(DWS)
- Oracle->PostgreSQL
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- Oracle->DDM
- DB2 for LUW -> GaussDB primary/standby
- DB2 for LUW -> GaussDB distributed
- DB2 for LUW->GaussDB(DWS)
- TiDB->GaussDB(for MySQL)
- Microsoft SQL Server->GaussDB(DWS)
- Microsoft SQL Server -> GaussDB primary/standby
- Microsoft SQL Server -> GaussDB Distributed
- Microsoft SQL Server->Microsoft SQL Server
- Microsoft SQL Server->PostgreSQL
- MongoDB->DDS
- MariaDB->MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- GaussDB(for MySQL)->GaussDB(for MySQL)
- Cassandra->GeminiDB Cassandra
- Dynamo->GeminiDB Dynamo

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- MySQL->MariaDB
- DDS->MongoDB
- DDS->Kafka
- DDM->MySQL
- DDM->Oracle
- DDM->Kafka
- GaussDB(for MySQL)->Kafka
- GaussDB(for MySQL)->CSS/ES
- GaussDB(for MySQL)->Oracle
- GaussDB(for MySQL)->MySQL
- GaussDB(for MySQL)->GaussDB(DWS)
- GaussDB primary/standby -> GaussDB primary/standby
- GaussDB primary/standby -> GaussDB distributed
- GaussDB primary/standby -> GaussDB(DWS)
- GaussDB primary/standby -> MySQL
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> Kafka
- GaussDB distributed -> GaussDB distributed
- GaussDB distributed -> GaussDB primary/standby
- GaussDB distributed -> GaussDB(DWS)
- GaussDB distributed -> MySQL
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->PostgreSQL
- PostgreSQL->Kafka
- MariaDB->MariaDB
- Microsoft SQL Server->Kafka

Self-built -> Self-built

- MySQL->Kafka
- MySQL->CSS/ES
- MySQL -> GaussDB Primary/Standby
- MySQL -> GaussDB Distributed
- Oracle-> Kafka
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed

- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> Kafka
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->Kafka
- DB2 for LUW -> GaussDB primary/standby
- DB2 for LUW -> GaussDB distributed
- Microsoft SQL Server->Kafka

7.7.11 Resetting Synchronization Position

For real-time synchronization tasks from MySQL to GaussDB(DWS), DRS allows you to reset the incremental synchronization position to place the synchronization task to the time before a fault occurs and perform synchronization again.

Constraints

- Only synchronization tasks from MySQL to GaussDB(DWS) support this function.
- When a task is in the incremental phase, you need to pause the task and then reset the synchronization position.

Procedure

Step 1 On the **Data Synchronization Management** page, locate the target synchronization task and choose **More > Reset Position** in the **Operation** column.

Figure 7-52 Reset Position

Task Name/ID	Status	Delay	Dist...	Del...	DB En...	Synchr...	Created	N...	Billing ...	Description	Enter...	Operation
DRS-4375- 990125ae-2d88a-4e48-8ba5-2827ce...	Incremental	0s	Yes	To the cl...	MySQL	Full-increm...	Oct 10, 2023 15:2...	VPC	Yearly/Monthly Created on ...		default	Edit Speed More
DRS-6767- 9a249c29-7556-4c21-b98f-d03a085a...	Failed	--	Yes	To the cl...	MySQL-GA...	Full-increm...	Oct 10, 2023 15:2...	VPC	Pay-per-Use Created on ...	2	default	Resume Reset More
DRS-6313- 4f193225-1762-43ca-95aa-e4c017f9...	Incremental	0s	Yes	To the cl...	MySQL-GA...	Full-increm...	Oct 07, 2023 16:3...	VPC	Yearly/Monthly Created on ...	22	default	Edit Reset Position

Step 2 In the **Reset Position** dialog box, specify **Position in Time** and click **Check**.

Step 3 After verifying the synchronization position, click **Start** to submit the synchronization task again.

----End

7.7.12 Restarting a Synchronization Task

DRS allows you to restart a synchronization task after task parameters are changed in **Changing Task Parameters**.

Constraints

- Tasks in the **Full, Full synchronization failed, Incremental, or Incremental synchronization failed** state can be restarted.
- Only the synchronization tasks of the following engines can be restarted:

- MySQL->MySQL
- MySQL->PostgreSQL
- MySQL -> GaussDB Distributed
- MySQL -> GaussDB Primary/Standby
- MySQL->GaussDB(DWS)
- MySQL->GaussDB(for MySQL)
- MySQL->Kafka
- MySQL->CSS/ES
- MySQL->Oracle

Method 1

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Restart** in the **Operation** column.

Step 2 In the displayed **Restart Task** dialog box, click **Yes**.

----End

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task in the **Task Name/ID** column.

Step 2 On the **Basic Information** page, click **Restart** in the upper right corner.

Step 3 In the displayed **Restart Task** dialog box, click **Yes**.

----End

7.7.13 Skipping DDL Statements

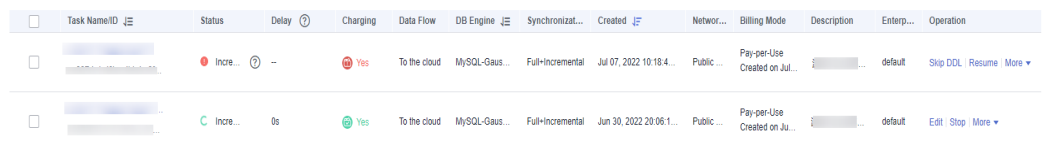
In a synchronization task, if the destination is a GaussDB(DWS) instance, DDL operations that are performed on the source database in the incremental synchronization phase will also be performed on the GaussDB(DWS) instance. If the destination GaussDB(DWS) instance does not support the DDL statement, the synchronization task becomes faulty.

To solve this problem, you can skip the DDL statement to continue the synchronization task after rectifying the fault based on the synchronization log information.

Procedure

Step 1 Contact GaussDB(DWS) technical support to execute this statement with the same semantics in the destination database and the statement must comply with the syntax of the destination database.

Step 2 On the **Data Synchronization Management** page, locate the task you want to skip and click **Skip DDL** in the **Operation** column.

Figure 7-53 Skip DDL

Task Name/ID	Status	Delay	Charging	Data Flow	DB Engine	Synchronization	Created	Network	Billing Mode	Description	Enterprise	Operation
[Redacted]	Incre.	--	Yes	To the cloud	MySQL-Gaus...	Full+Incremental	Jul 07, 2022 10:18:4...	Public	Pay-per-Use Created on Jul...	[Redacted]	default	Skip DDL, Resume, More
[Redacted]	Incre...	0s	Yes	To the cloud	MySQL-Gaus...	Full+Incremental	Jun 30, 2022 20:06:1...	Public	Pay-per-Use Created on Ju...	[Redacted]	default	Edit, Stop, More

Step 3 After confirming that statement was executed on the destination database, click **Yes** to skip the error and continue the synchronization task.

----End

7.7.14 Performing a Switchover for a Dual-AZ Task

You can set **DRS Task Type** to **Single-AZ** or **Dual-AZ** when creating a DRS real-time synchronization task. The dual-AZ deployment provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization.

You can select the DRS task type in the following scenarios:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->MariaDB
 - GaussDB(for MySQL)->GaussDB(for MySQL)
 - MariaDB->MySQL
 - MariaDB->GaussDB(for MySQL)
 - Oracle -> GaussDB Primary/Standby
 - Oracle -> GaussDB Distributed
 - Microsoft SQL Server->MySQL
 - Microsoft SQL Server->GaussDB(for MySQL)
 - Microsoft SQL Server->PostgreSQL
- From the cloud
 - MySQL->Kafka
 - MySQL->MariaDB
 - GaussDB(for MySQL)->Kafka
 - GaussDB(for MySQL)->CSS/ES
 - GaussDB primary/standby -> Kafka
 - GaussDB Distributed -> Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - Oracle->Kafka
 - Oracle -> GaussDB Primary/Standby
 - Oracle -> GaussDB Distributed

- GaussDB primary/standby -> Kafka
- GaussDB Distributed -> Kafka

Prerequisites

- You have logged in to the DRS console.

Scenarios

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, configure the task name, description, and the synchronization instance details, set **DRS Task Type** to **Dual-AZ**, and click **Next**.
- Step 3** Return to the **Data Synchronization Management** page, you can find that there are two subtasks displayed under the synchronization task you created.

Figure 7-54 Primary/Standby tasks

Task Name/ID	Status	Delay	Changing	Data Flow	DB Engine	Synchronization	Created	Network	Billing Mode	Description	Operation
DRS-6857-test 69611a04-0200-4c05-8aac-f...	Creating	--	No	To the cloud	MySQL	Full-incremental	May 31, 2022 15:08:09...	Public n...	Pay-per-use	--	Edit Stop Speed
DRS-6857-test-ch16-01 6962943-e50e-4960-a0b2-...	Creat...	--	No	To the clo...	MySQL	Full-incremen...	May 31, 2022 15:08:1...	Public ...	Pay-per-use	--	
DRS-6857-test-ch16-02 69626ac-5022-450c-a8ff-f...	Creat...	--	No	To the clo...	MySQL	Full-incremen...	May 31, 2022 15:08:1...	Public ...	Pay-per-use	--	

- Step 4** After a synchronization task is configured and started, DRS will start the task in the primary AZ, and the task in the standby AZ is suspended.

Figure 7-55 Before the primary/standby switchover

Task Name/ID	Status	Delay	Changing	Data Flow	DB Engine	Synchronization	Created	Network	Billing Mode	Description	Operation
DRS-6855-test 4023c78c-594c-4904-8ba0-f...	Increm...	0s	Yes	To the cloud	MySQL	Full-incremental	May 31, 2022 10:03:08...	Public n...	Pay-per-use	--	Edit Stop More
DRS-6855-ch16-01 706c25f-0943-4a24-a104-...	Incre...	0s	Yes	To the clo...	MySQL	Full-incremen...	May 31, 2022 10:03:1...	Public ...	Pay-per-use	--	
DRS-6855-ch16-02 775315d8-d8ff-4968-81aa-...	Paused	--	Yes	To the clo...	MySQL	Full-incremen...	May 31, 2022 10:03:1...	Public ...	Pay-per-use	--	

- Step 5** If the task in the primary AZ is abnormal, DRS automatically starts the task in the standby AZ to continue the synchronization.

Figure 7-56 After the primary/standby switchover

Task Name/ID	Status	Delay	Changing	Data Flow	DB Engine	Synchronization	Created	Network	Billing Mode	Description	Operation
DRS-6855-test 4023c78c-594c-4904-8ba0-f...	Increm...	0s	Yes	To the cloud	MySQL	Full-incremental	May 31, 2022 10:03:08...	Public n...	Pay-per-use	--	Edit Stop More
DRS-6855-ch16-01 706c25f-0943-4a24-a104-...	Paused	--	Yes	To the clo...	MySQL	Full-incremen...	May 31, 2022 10:03:1...	Public ...	Pay-per-use	--	
DRS-6855-ch16-02 775315d8-d8ff-4968-81aa-...	Incre...	0s	Yes	To the clo...	MySQL	Full-incremen...	May 31, 2022 10:03:1...	Public ...	Pay-per-use	--	

----End

7.7.15 Exchanging the Direction for a Two-Way Synchronization Task

In two-way synchronization, only forward tasks support DDL execution to prevent DDL loopback. DRS allows you to exchange the direction of a synchronization task. You can use this function to change the task role to enable DDL execution on backward tasks.

Constraints

- Only two-way synchronization tasks support this function.
- The direction can be exchanged only when both the forward and backward tasks are paused.
- You need to resume the task to apply the change.

Procedure

- Step 1** On the **Data Synchronization Management** page, locate the paused two-way synchronization task.
- Subtask 1 is a forward task.
- Step 2** Click **Exchange Direction** in the **Operation** column of the task.
- Step 3** In the displayed dialog box, click **Yes**.
- Step 4** After the direction exchange, view that the synchronization relationship of subtask 1 changes and subtask 1 becomes a backward task.
- Step 5** Click **Resume** in the **Operation** column of the subtask.

----End


7.7.16 Changing a Single-AZ Task to a Dual-AZ Task

DRS allows you to change a single-AZ task to a dual-AZ task, improving task reliability while remaining the original task.

Constraints

- A dual-AZ task cannot be changed to a single-AZ task.
- Only tasks in the **Incremental**, **Incremental failed**, or **Paused** state can be changed.
- Only synchronization tasks from GaussDB Primary/Standby to Kafka can be changed.

Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** On the **Data Synchronization Management** page, select the target task and choose **More > Change to Primary/Standby** in the **Operation** column.
- Step 5** On the displayed page, select the standby AZ. If the task is over a public network, you need to specify the EIP of the standby task and click **OK**.
- Step 6** After submitting the change, click **Back to Task List**. On the **Data Synchronization Management** page, the task status is **Changing to primary/standby**.

After the task type change is complete, you can check that the parent task has two subtasks: one primary task and one standby task.

----End

7.7.17 Cloning a Synchronization Task

DRS allows you to clone the configuration of existing synchronization tasks. However, tasks in the following status cannot be cloned:

- Creating
- Creation failed
- Configuration
- Pending start
- Starting
- Deleted

You can clone the following data flow types:

- To the cloud
 - MySQL->MySQL
 - MySQL->MariaDB
 - PostgreSQL->PostgreSQL
 - Microsoft SQL Server->MySQL
 - Microsoft SQL Server->GaussDB(for MySQL)
 - Microsoft SQL Server->PostgreSQL
- From the cloud
 - MySQL->MySQL
 - MySQL->CSS/ES
 - MySQL->MariaDB
 - DDS->MongoDB
- Self-built -> Self-built
 - MySQL->CSS/ES

NOTE

- When a task is cloned, the source and destination database passwords are not cloned. You need to enter the passwords again for the new task.
- When you clone a task, the advanced settings for data filtering are not cloned. You need to set the advanced settings for the cloned task again.
- Many-to-one task cloning is not supported.
- When you clone a task that is being changed, if the change information has been saved to the database, the clone task configuration is the same as the changed configuration.
- After a clone task is created, the EIP and private IP address of the new task are different from those of the original task. You may need to configure the network to ensure that the new task can communicate with the source and destination databases.

Prerequisites

- You have logged in to the DRS console.

- A synchronization task has been created.

Procedure

- Step 1** On the **Data Synchronization Management** page, select the task to be cloned and click **Clone** in the **Operation** column.
- Step 2** In the displayed dialog box, confirm the new task name and click **OK**.
- Step 3** After the task is submitted and the task clone is complete, the task status changes to **Configuration**. You can click **Edit** in the **Operation** column, enter the source and destination database passwords again, and edit and start the task.

----End


7.7.18 Changing Specifications

You can change the DRS task specifications based on your service requirements. After the specification change starts, the task enters the **Changing specifications** state and data synchronization is suspended. After the specification change is complete, the task is automatically resumed.

Constraints

- You can change the task specifications only when your account balance is more than \$0 USD.
- DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be upgraded for dual-AZ tasks or downgraded.
- DRS allows you to change the specifications of synchronization tasks only in the **Full**, **Full failed**, **Incremental**, or **Incremental failed** state.
- Before changing the task specifications, ensure that the current AZ supports the target specifications.
- You are advised to change the task specifications during off-peak hours.
- After the specification change starts, the task is suspended. The task is automatically resumed after the change is complete.
- It takes about 5 to 10 minutes to change the task specifications.

Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** On the **Data Synchronization Management** page, select the target task and choose **More > Change Specifications** in the **Operation** column.
- Step 5** On the displayed page, select the desired specifications, perform a pre-check, and click **Next**.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances, click **Change**.
To view the cost incurred by the specifications change, choose **Billing Center > Cost Bills** in the upper right corner.
- For yearly/monthly DB instances, click **Change**. On the displayed page, click **Pay**. You can change the specifications only after the payment is successful.

Step 7 View the task specification change result.

After the application is submitted, click **Back to Task List**. On the **Data Synchronization Management** page, the instance status is **Changing specifications**.

After the task status changes from **Changing specifications** to another status, you can view the instance specifications on the **Basic Information** page to check whether the change is successful. Alternatively, you can view the change logs on the **Synchronization Logs** page to whether the change is successful.

- **change specification start**: indicates that the specification change starts.
- **change specification success**: indicates that the specifications are changed.
- **change specification failed**: indicates that the specifications fail to be changed.

Figure 7-57 Specification change logs

Time	Level	Description
2022/12/02 19:24:05 GMT+08:00	Info	change specification failed
2022/12/02 19:15:33 GMT+08:00	Info	pause job complete
2022/12/02 19:15:28 GMT+08:00	Info	start to pause the current job
2022/12/02 19:14:53 GMT+08:00	Info	precheck [acecaac5-782c-4a9d-9344-d5cb221j0201] completed, total item:4, success item:4, not pass item:0
2022/12/02 19:14:32 GMT+08:00	Info	precheck [acecaac5-782c-4a9d-9344-d5cb221j0201] start
2022/12/02 19:14:18 GMT+08:00	Info	precheck [acecaac5-782c-4a9d-9344-d5cb221j0201] completed, total item:4, success item:4, not pass item:0
2022/12/02 19:13:57 GMT+08:00	Info	precheck [acecaac5-782c-4a9d-9344-d5cb221j0201] start
2022/12/02 19:12:45 GMT+08:00	Info	change specification start
2022/12/02 18:45:17 GMT+08:00	Error	service LOGMANAGER failed, cause by: database log download failed, error code is 'code': '01300', 'name': 'LOGS_NOT_EXIST', 'retry': false, 'reset': false, 'level': 3
2022/12/02 18:44:52 GMT+08:00	Info	increment transfer start

----End

7.7.19 Unsubscribing from a Yearly/Monthly Task


To delete a DRS task billed on the yearly/monthly basis, you need to unsubscribe the order.

Prerequisites

- You have logged in to the DRS console.
- The billing mode of the current DRS instance is yearly/monthly.

Method 1

Unsubscribe from a yearly/monthly task on the **Data Synchronization Management** page.

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** On the **Data Synchronization Management** page, select the target task and choose **More > Unsubscribe** in the **Operation** column.
- Step 5** In the displayed dialog box, click **Yes**. The **Unsubscribe from Resource** page is displayed.
- Step 6** On the **Unsubscribe from Resource** page, verify the information about the instance to be unsubscribed, specify a reason, select the check box, and click **Confirm**.

 **NOTE**

After a DRS instance is unsubscribed, the DRS task ends immediately. Ensure that data synchronization is complete or the DRS instance is no longer used.

- Step 7** In the displayed dialog box, click **Yes**.

----End

Method 2

Unsubscribe from a yearly/monthly task on the **Billing Center** page.


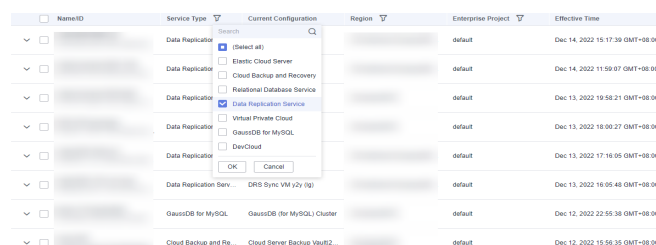
- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** Click **Billing & Costs** from the top menu bar. The **Billing Center** page is displayed.
- Step 5** In the navigation pane, choose **Orders > Unsubscriptions**.
- Step 6** On the displayed page, select the order to be unsubscribed and click **Unsubscribe** in the **Operation** column.
 - You can select DRS in the **Service Type** to filter all DRS orders.

Figure 7-58 Filtering all orders



NameID	Service Type	Current Configuration	Region	Enterprise Project	Effective Time
	Data Replication	(Select all)		default	Dec 14, 2022 15:17:39 GMT+08:00
	Data Replication	<input type="checkbox"/> Elastic Cloud Server		default	Dec 14, 2022 11:59:07 GMT+08:00
	Data Replication	<input type="checkbox"/> Cloud Backup and Recovery		default	Dec 13, 2022 19:58:21 GMT+08:00
	Data Replication	<input checked="" type="checkbox"/> Data Replication Service		default	Dec 13, 2022 18:00:27 GMT+08:00
	Data Replication	<input type="checkbox"/> Virtual Private Cloud		default	Dec 13, 2022 17:18:05 GMT+08:00
	Data Replication	<input type="checkbox"/> GaussDB for MySQL		default	Dec 13, 2022 17:18:05 GMT+08:00
	Data Replication	<input type="checkbox"/> DevCloud		default	Dec 13, 2022 16:05:48 GMT+08:00
	Data Replication Serv...	DRS Sync VM (2y) (9)		default	Dec 12, 2022 22:55:38 GMT+08:00
	GaussDB for MySQL	GaussDB (for MySQL) Cluster		default	Dec 12, 2022 15:56:35 GMT+08:00
	Cloud Backup and Re...	Cloud Server Backup Vault2...		default	Dec 12, 2022 15:56:35 GMT+08:00

- Alternatively, search for target orders by name, order No., or ID in the search box.

Step 7 On the displayed page, confirm the order to be unsubscribed from and select a reason. Then, click **Confirm**.

For unsubscription details, see [Unsubscription Rules](#).

Step 8 In the displayed dialog box, click **Yes**.

 **NOTE**

After a DRS instance is unsubscribed, the DRS task ends immediately. Ensure that data synchronization is complete or the DRS instance is no longer used.

----End

7.7.20 Stopping a Synchronization Task

After the source database and services are migrated to the destination database, you can stop the synchronization task. To prevent data from being overwritten after the source database and services are migrated to the destination database, stop a synchronization task to achieve this goal.

You can stop a task in any of the following statuses:

- Creating
- Configuration
- Pending start
- Full synchronization
- Full synchronization failed
- Incremental synchronization
- Incremental synchronization failed
- Paused
- Fault rectification

NOTICE

- You are advised to stop the task before performing other operations, such as disconnecting the network between the source database and the synchronization instance. Otherwise, an alarm indicating that the source database cannot be connected will be generated.
 - For a task in the **Configuration** state, it cannot be stopped if it fails to be configured.
 - For a task in the **Fault rectification** state, it cannot be stopped if the fault is being rectified.
 - After a task is stopped, it cannot be retried.
-

Procedure

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Stop**.

Step 2 In the displayed dialog box, click **Yes**.

NOTE

- If the task status is abnormal (for example, the task fails or the network is abnormal), DRS will select **Forcibly stop task** to preferentially stop the task to reduce the waiting time.
- Forcibly stopping a task will release DRS resources. Check whether the synchronization is affected.
- To stop the task properly, restore the DRS task first. After the task status becomes normal, click **Stop**.
- For a DRS task that is in the incremental state and with MySQL serving as the source database, after you select **Display breakpoint information when the task is stopped** when you stop the task, the GTID and binlog position information of the source database will be displayed on the synchronization progress page after the task is stopped.


----End

Stopping Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be stopped.

Step 2 Click **Batch Operations** in the upper left corner and choose **Stop**.

Figure 7-59 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

7.7.21 Deleting a Synchronization Task

This section describes how to delete a synchronization task that has been completed or has failed. Deleted tasks will no longer be displayed in the task list. Exercise caution when performing this operation.

Prerequisites

You have logged in to the DRS console.

Deleting a Task

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Delete** in the **Operation** column.

Step 2 Click **Yes** to submit the deletion task.


----End

Deleting Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be deleted.

Step 2 Click **Batch Operations** in the upper left corner and choose **Delete**.

Figure 7-60 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

7.7.22 Importing Tasks in Batches

DRS allows you to import an Excel file to create synchronization tasks in batches. You need to fill in a template with basic task information, source and destination database information, and objects to be synchronized, and then import the template.

Constraints

- Only real-time synchronization tasks from MySQL to MySQL, from MySQL to GaussDB Primary/Standby, from MySQL to GaussDB Distributed, from Oracle to GaussDB Primary/Standby, from Oracle to GaussDB Distributed, from GaussDB Primary/Standby to GaussDB Primary/Standby, from GaussDB Primary/Standby to GaussDB Distributed, from GaussDB Primary/Standby to Kafka, from GaussDB Distributed to GaussDB Distributed, from GaussDB Distributed to GaussDB Primary/Standby, and from GaussDB Distributed to Kafka can be imported.
- When editing the template for importing tasks, clear the example content and enter the required values.
- When specifying the source and destination databases, you can set only mandatory parameters. You need to ensure that the subnets of the source and destination databases can communicate with each other.
- The import template is designed based on the API and GUI information. You need to set the parameters based on the description and ensure the parameter accuracy.
- After the import template is uploaded, the system applies for creating DRS instances. Some parameters cannot be modified. Exercise caution when performing this operation.
- After the import template is uploaded, the system automatically saves all parameters for the task and performs steps such as connection test, object selection, and pre-check. After the pre-check is successful, the task is in the **Configuration** state. You can click **Start Now** to start the task, or click **Edit** to modify other parameters.
- You cannot schedule the start time of a task in the import template. After the pre-check for the task is successful, you can click **Edit** on the GUI to access the startup page and set a scheduled start time.

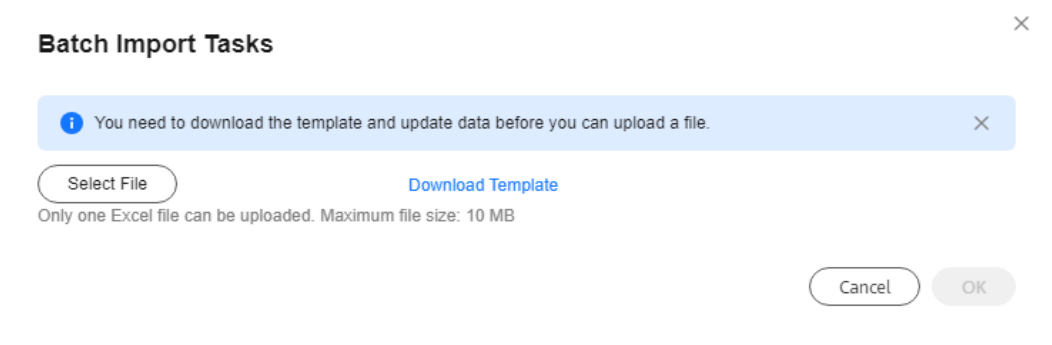
- The system cannot verify the correctness of all parameters in the import template. You are advised to click **Edit** to enter the page, confirm the correctness of the data saved in each step, and then start the task.
- The import template is a common template. One sheet is used to create one task. To create tasks in batches, copy multiple sheets and set parameters.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Batch Import Tasks** in the upper left corner.

Step 2 On the **Batch Import Tasks** dialog box, click **Download Template**.

Figure 7-61 Batch Import Tasks



Step 3 Edit the downloaded Excel template and enter task information as required.

Step 4 Click **Select File** to upload the template.

Step 5 After the upload and verification are successful, click **Confirm** to create tasks in batches.

NOTE

- DRS does not save the uploaded Excel file and only parses the content in the file.
- The IP addresses, domain names, usernames, and passwords of the source and destination databases are encrypted and stored in the system until a task is deleted.

----End

7.7.23 Filtering Large Fields

DRS allows you to filter large fields by setting a field filtering threshold and replacement characters to process special fields (blob, mediumblob, longblob, varbinary, mediumtext and longtext) in a synchronization table.

Constraints


- Large field filtering is supported only for tasks from MySQL to Kafka and from GaussDB(for MySQL) to Kafka.
- When creating a task, you can choose to enable or disable large field filtering in object selection page. After a task is started, you can choose **More > Filter Large Field** in the **Operation** column of the task to set this function.

- For a failed task, you need to manually resume or reset the task to apply this function.

Procedure

Step 1 On the **Data Synchronization Management** page, locate the target task and choose **More > Filter Large Field** in the **Operation** column.

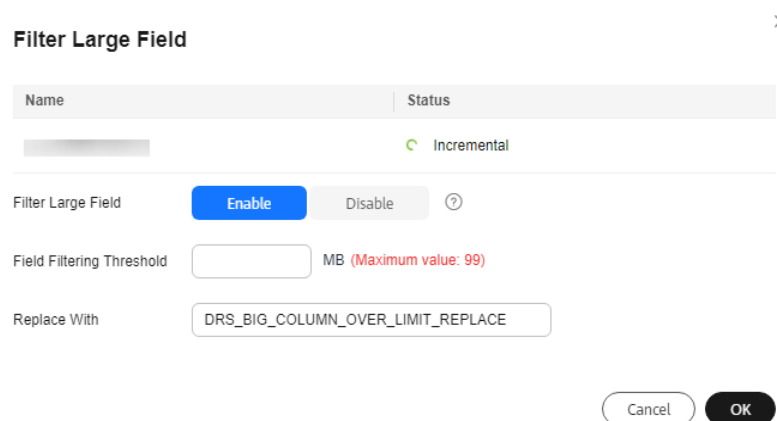
Figure 7-62 Task list



Step 2 In the **Filter Large Field** dialog box, choose whether to use large field filtering to process special fields (blob, mediumblob, longblob, varbinary, mediumtext and longtext) in a synchronization table.

- **Enable:** You need to set **Field Filtering Threshold** and **Replace With**. If the size of a field exceeds the threshold, the value is replaced based on a specified character. Note that large field filtering is used to replace the value of a field, not the entire DML record. If a DML record contains many large fields, the size of only some of these fields exceeds the filtering threshold, and the accumulated value of other fields that do not exceed the filtering threshold is greater than the value of **request.max.size**, when data is written to Kafka, the size of the message body in the destination Kafka database may still exceed the upper limit, resulting in a DRS error.
- **Disable:** Large fields are not filtered.

Figure 7-63 Filter Large Field



Step 3 Click **OK**.

For a failed task, you need to manually resume or reset the task to apply this function.

----End

7.7.24 Upgrading the Version of a DRS Task

In 24.10.0 and later versions, you can upgrade the version of a DRS task on the console immediately or as scheduled. You can set an upgrade time window. During the time window, the system checks whether a task meets the upgrade conditions every 10 minutes. If the task meets the upgrade conditions, the system delivers the request to the kernel for version upgrade.

Constraints

- The upgrade conditions displayed on the console are as follows:
 - The task is in the incremental state.
 - The kernel version must be 24.10.0 or later. To upgrade the DRS kernel, submit a service ticket by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.
 - The task has a version that can be upgraded.
 - The version upgrade function is available only for synchronization, migration, and DR tasks.
- The upgrade conditions on which the management system depends for delivering an upgrade command are as follows:
 - The task is in the incremental state.
 - The task latency is no more than 30s.
 - The task has a version that can be upgraded, and the upgrade time is within a specified time range.
 - In the multi-task scenario, an upgrade can be performed only when all subtasks meet the preceding conditions.
- After receiving the upgrade request from the management system, the kernel downloads the new version and monitors and checks whether the upgrade is successful based on the following conditions:
 - The log download, parsing, and incremental migration processes are normal.
 - The incremental position is updated properly.
- The maximum monitoring duration is 10 minutes. (If the system detects that the incremental position is updated properly within 3 minutes and the preceding processes are normal within 3 minutes, the monitoring terminates within 3 minutes.)
- If the task upgrade fails, the system rolls back the version. The entire upgrade process is displayed in the **Migration Logs** page on the console.
- The task cannot be paused during version upgrade.

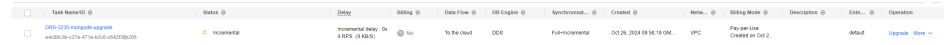
Prerequisites

- The task is in the incremental state. The task has been upgraded to the baseline version. The **Upgrade** button is available in the task list.

Procedure

Step 1 On the **Data Synchronization Management** page, locate the target task and click **Upgrade** in the **Operation** column.

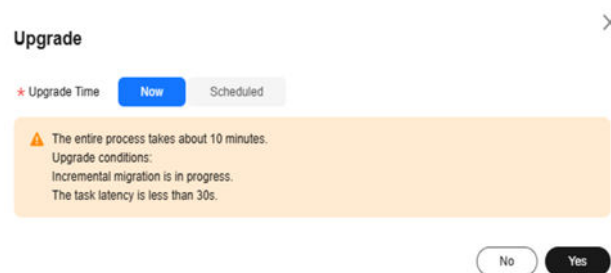
Figure 7-64 Task management



Step 2 In the **Upgrade** dialog box, select **Now** or **Scheduled** for **Upgrade Time**.

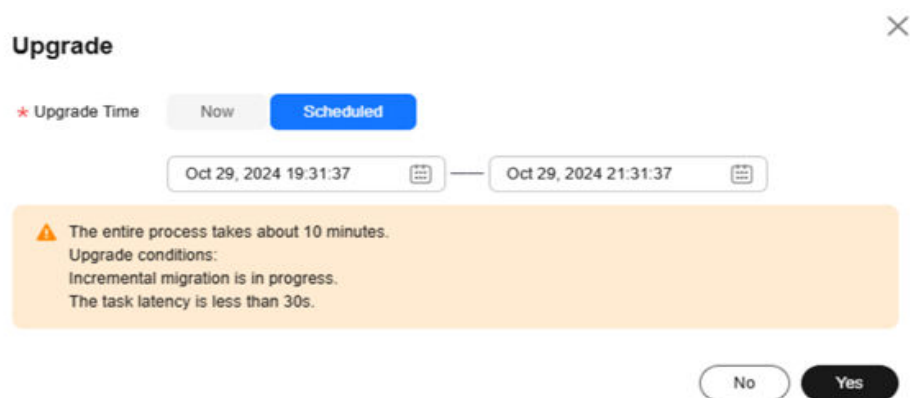
1. After you click **Now**, the task upgrade starts.

Figure 7-65 Upgrading a task



2. If you select **Scheduled** for **Upgrade Time**, specify a time range for task upgrade.

Figure 7-66 Upgrading a task as scheduled



----End

Viewing the Upgrade Status

- During the upgrade, the task status is **Upgrading replication instance** in the task list.

Figure 7-67 Upgrading



- Click the task name to go to the task details page. In the navigation pane, choose **Synchronization Logs** to view upgrade logs.
 - Upgrade success logs

Figure 7-68 Logs

Time	Level	Description
Oct 26, 2024 12:10:13 GMT+08:00	info	version upgrade from 24.10.0.2 to 24.10.0.4 completed, status successful
Oct 26, 2024 12:14:42 GMT+08:00	info	increment transfer start
Oct 26, 2024 12:14:16 GMT+08:00	info	version upgrade from 24.10.0.2 to 24.10.0.4 started
Oct 26, 2024 12:12:27 GMT+08:00	info	increment transfer start
Oct 26, 2024 12:12:19 GMT+08:00	info	pause job complete

- Upgrade failure logs

Figure 7-69 Logs

Time	Level	Description
Oct 26, 2024 12:00:03 GMT+08:00	info	version upgrade from 24.10.0.0 to 24.10.0.1 completed, status failed
Oct 26, 2024 12:06:03 GMT+08:00	info	rollback to version 24.10.0.0 successful
Oct 26, 2024 12:05:52 GMT+08:00	info	increment transfer start
Oct 26, 2024 12:05:38 GMT+08:00	info	version upgrade from 24.10.0.0 to 24.10.0.1 failed, errMsg:task:CheckUpgradeStatusTask error, rollback to version 24.10.0.0 started
Oct 26, 2024 12:04:08 GMT+08:00	info	increment transfer start
Oct 26, 2024 12:03:27 GMT+08:00	info	version upgrade from 24.10.0.0 to 24.10.0.1 started

7.7.25 Task Statuses

Synchronization statuses indicate different synchronization phases.

Table 7-14 lists synchronization task statuses and descriptions.

Table 7-14 Task status description

Status	Description
Creating	A synchronization instance is being created.
Task creation failed	Failed to create a real-time synchronization instance.
Configuration	The synchronization instance is successfully created, but the synchronization task is not started. You can continue to configure the task.
Frozen	Instances are frozen when the account balance is less than or equal to \$0.
Pending start	The scheduled synchronization task has been delivered to the synchronization instance, waiting for the synchronization instance to start the synchronization task.
Starting	The task is being started.

Status	Description
Start failed	A real-time synchronization task fails to be started.
Full synchronization	A full synchronization task is being performed.
Full synchronization failed	A full synchronization task fails.
Incremental synchronization	An incremental synchronization task is being performed.
Incremental synchronization failed	An incremental synchronization task fails.
Modifying task	The synchronization object is being modified.
Modifying task failed	The synchronization object fails to be modified.
Fault rectification	A synchronization instance is faulty and the system automatically restores the synchronization task.
Paused	The real-time synchronization task has been paused.
Cloning	A synchronization task is being cloned.
Cloning failed.	The clone synchronization task fails.
Changing specifications	The instance specifications are being changed. After the change is complete, the task is automatically restored.
Specification change failed	The instance specifications fail to be changed.
AZ switchover failed	The AZ switchover fails for the real-time synchronization task.
Task stopping	The synchronization instance and resources used for executing the synchronization task are being released.
Completing	A synchronization instance and resources are being released.
Stopping task failed	The synchronization instance and resources fail to be released.
Completed	The task is completed and the synchronization instance is released.

 **NOTE**

- If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources again.
- Deleted synchronization tasks are not displayed in the status list.

8 Tag Management

Scenarios

Tag Management Service (TMS) enables you to use tags on the management console to manage resources. TMS works with other cloud services to manage tags. TMS manages tags globally, and other cloud services manage their own tags. If you have to manage a large number of tasks, you can use different tags to identify and search for tasks.


- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each DB instance can have up to 20 tags.

Adding a Tag

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, click **Edit Tag**. In the displayed dialog box, click **Add Tag**, enter a tag key and value, and click **OK**.



- When you enter a tag key and value, the system automatically displays all tags (including predefined tags and resource tags) associated with all DB instances except the current one.

- The tag key cannot be empty and must be unique. It cannot start or end with a space or start with **_sys_**. It can contain 1 to 128 characters, including letters, digits, spaces, and special characters **_:=+@**
- The tag value can be empty. It cannot start or end with a space and can contain 0 to 255 characters, including letters, digits, spaces, and special characters **_:=+@**

Step 4 View and manage the tag on the **Tags** page.

----End

Editing a Tag

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, click **Add/Edit Tags**. In the displayed dialog box, modify the tag and click **OK**.

----End

Delete a Tag

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

Step 4 After the tag is deleted, it will no longer be displayed on the **Tags** page.

----End

9 Connection Diagnosis

If a DRS instance fails to be connected to the source or destination database during connection testing, DRS provides the quick diagnosis function and returns the diagnosis result.

- You can perform connection diagnosis only on the task node whose database information is obtained by entering an IP address or selecting a task node on the GUI. DN diagnosis of GaussDB is not supported.
- In cluster or multi-AZ task scenarios, diagnosis can be performed only on the node of the primary task.

Prerequisites

- You have logged in to the DRS console.
- A task has been created.

Procedure

- Step 1** On the task management page, click the target task name in the **Task Name/ID** column.
- Step 2** On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the DRS instance.

If the connection testing fails, click **Quick Diagnosis** on the right of the failure information to diagnose the fault.

Figure 9-1 Quick Diagnosis

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Test Connection

The network connection between the replication instance and the database is faulty. For a cross-VPC task, the network between different VPCs may be disconnected. For details about how to create a VPC peering connection, see the VPC documentation, [View details](#)

Quick Diagnosis

Step 3 View the diagnosis result on the displayed **Diagnosis Details** dialog box. The result includes the packet loss rate and port check result.

Figure 9-2 Diagnosis Details

Diagnosis Details ×

IP Address or Domain Name	Packet Loss Rate (%)	Port Check
<input type="text"/>	100	● Failed

OK

----End

10 Viewing Abnormal Data

During the synchronization, data may become abnormal due to conflicts, data processing, heterogeneous conversion, and object missing. You can view abnormal data on the **Abnormal Records** tab to locate the fault.

 **NOTE**

This function is available for data synchronization from DDM to GaussDB(DWS), DDM to Oracle, Oracle to GaussDB(DWS), Oracle to RDS for MySQL, Oracle to GaussDB(for MySQL), Oracle to PostgreSQL, and MySQL to GaussDB(DWS).

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the **Abnormal Records** tab, filter abnormal data, including statements and description, by time range.

----End

11 Interconnecting with CTS

11.1 Key Operations Recorded by CTS

Cloud Trace Service (CTS) provides records of operations on cloud service resources, enabling you to query, audit, and backtrack operations.

Table 11-1 DRS operations recorded by CTS

Operation	Resource Type	Trace Name
Creating a task	job	createJob
Editing a task	job	modifyJob
Deleting a task	job	deleteJob
Starting a task	job	startJob
Resuming a task	job	retryJob

11.2 Viewing Traces

After CTS is enabled, CTS starts recording operations on cloud resources. The CTS management console stores the last seven days of operation records.



This section describes how to query the operation records of the last seven days on the CTS console.

Prerequisites

The CTS service has been enabled.

Procedure

Step 1 Log in to the management console.

- Step 2** Click  in the upper left corner of the page and select a region and project.
- Step 3** Click **Service List**. Under **Management & Governance**, choose **Cloud Trace Service**.
- Step 4** Choose **Trace List** in the navigation pane on the left.
- Step 5** Specify the search criteria as needed.
- **Search time range:** In the upper right corner, choose **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range.
 - **Trace Type, Trace Source, Resource Type, and Search By:** Select a filter from the drop-down list.
If you select **Resource ID** for **Search By**, specify a resource ID.
If you select **Data** for **Trace Type**, you can only filter traces by tracker.
 - **Operator:** Select a specific operator (a user rather than a tenant).
 - **Trace Status:** Available options include **All trace statuses**, **normal**, **warning**, and **incident**. You can only select one of them.
- Step 6** Click **Query**.
- Step 7** Click  to the left of the target record to extend its details.
- Step 8** Click **View Trace** in the **Operation** column. A dialog box is displayed, on which the trace structure details are displayed.

----End

12 Interconnecting with Cloud Eye

12.1 Supported Metrics

Description

This section describes metrics reported by the Data Replication Service (DRS) to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for DRS.

Namespace

SYS.DRS

DB Instance Monitoring Metrics

[Table 12-1](#) lists the DRS performance metrics.

Table 12-1 DRS metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_util	CPU Usage	CPU usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
mem_util	Memory Usage	Memory usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_incoming_bytes_rate	Network Input Throughput	Incoming traffic in bytes per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_outgoing_bytes_rate	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_read_bytes_rate	Disk Read Throughput	Number of bytes read from the disk per second (bytes/second).	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_write_bytes_rate	Disk Write Throughput	Number of bytes written to the disk per second (bytes/second).	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_util	Storage Space Usage	Storage space usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_bytes_rate	Source Database Read Throughput	Table data or WAL bytes read from the source database per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_rows_rate	Rows Read from Source Database per Second	Number of table data rows or WAL rows read from the source database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_latency	Source Database WAL Extract Lag	Latency of extracting WAL from the source database Unit: ms.	\geq ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_bytes_rate	Destination Database Write Throughput	Number of bytes written to the destination database per second.	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_rows_rate	Rows Written into Destination Database per Second	Number of rows that are written to the destination database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_transactions_rate	DML TPS	Number of DML transactions written to the destination database per second.	≥ 0 transaction/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_ddls_number apply_ddls_rate NOTE apply_ddls_rate is replaced by apply_ddls_number after December 2022.	DDL TPS	Total number of DDL transactions written into the destination database.	≥ 0 transaction	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_latency	Replication Delay	Delay (in milliseconds) of data replay.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_execute_time	Average Transaction Execution Time	Average execution time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is millisecond.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_commit_time	Average Transaction Commit Time	Average commit time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is ms.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_current_state	Synchronization Status	This metric is the synchronization status of the current kernel data (10: abnormal; 1: idle; 2: DML; 3: DDL), instead of the task status.	10: abnormal 1: idle 2: DML is executed. 3: DDL is executed.	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_thread_workers	Synchronization Threads	Number of working threads for data synchronization	≥ 0	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_job_status	Task Status	Status of the current task. (0: normal; 1: abnormal; 2: paused)	0: normal 1: abnormal 2: paused	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Dimensions

Key	Value
instance_id	DRS instance ID


12.2 Configuring Alarm Rules

Scenarios

You can configure DRS alarm rules to customize the monitored objects and notification policies and learn the DRS running status in a timely manner.

This section describes how to set DRS alarm rules, including the alarm rule name, service, dimension, monitoring scope, template, and whether to send a notification.

Procedure

- Step 1** Log in to the management console.
- Step 2** Under **Management & Governance**, click **Cloud Eye**.
- Step 3** In the navigation pane on the left, choose **Cloud Eye > Data Replication Service**.
- Step 4** Select the DB instance which you want to create an alarm rule for and click **Create Alarm Rule** in the **Operation** column.
- Step 5** On the displayed page, set parameters as required.
- Specify **Name** and **Description**.
 - Select **Use template** for **Method**. The template contains the following common metrics: CPU usage, memory usage, and storage space usage.
 - Click  to enable alarm notification. The validity period is 24 hours by default. If the topics you required are not displayed in the drop-down list, click **Create an SMN topic**. Then, select **Generated alarm** and **Cleared alarm** for **Trigger Condition**.

NOTE

Cloud Eye sends notifications only within the validity period specified in the alarm rule.

- Step 6** Click **Create**. The alarm rule is created.

For details about how to create alarm rules, see [Creating an Alarm Rule](#) in the *Cloud Eye User Guide*.

----End

12.3 Viewing Monitoring Metrics

Scenarios


Cloud Eye monitors the running statuses of replication, synchronization, and DR instances. You can obtain the monitoring metrics on the management console. Monitored data requires a period of time for transmission and display. The status of the monitored object displayed on the Cloud Eye page is the status obtained 5 to 10 minutes before. You can view the monitored data of a newly created DB instance 5 to 10 minutes later.

Prerequisites

An instance is running properly when in the following statuses:

- Real-time migration: Full migration and Incremental migration
- Real-time synchronization: Full synchronization and Incremental synchronization
- Real-time disaster recovery: Disaster recovery in progress


Viewing Metrics

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 5** On the displayed page, click **View Metric** in the upper right corner of the page to go to the Cloud Eye console.

By default, the monitoring information about the DRS instance is displayed on this page.

- Step 6** View monitoring metrics of the instance.
 - On the Cloud Eye console, click the target DB instance name and click **Select Metric** in the upper right corner. In the displayed dialog box, you can select the metrics to be displayed and sort them by dragging them at desired locations.
 - You can sort graphs by dragging them based on service requirements.
 - Cloud Eye can monitor performance metrics from the last 1 hour, 3 hours, 12 hours, 1 day, 7 days, and 6 months.

Figure 12-1 Viewing monitoring metrics



----End

13 Interconnecting with LTS

13.1 Log Reporting

Scenarios

If you enable log reporting, all logs generated by DRS instances (including real-time migration, backup migration, real-time synchronization, real-time disaster recovery, and workload replay instances) are uploaded to Log Tank Service (LTS) for management.

Precautions

- After this function is enabled, all logs of the task are reported by default.
- This request does not take effect immediately. There is a delay of about 10 minutes.
- You will be billed for this function. For details, see [LTS Pricing Details](#).
- Ensure that there are available LTS log groups and log streams in the same region as your instance.

For more information about log groups and log streams, see [Log Management](#).

- After this function is disabled, you will not be billed anymore.

Enabling or Disabling Log Reporting

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.

Step 4 Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column. The operations for real-time synchronization, real-time disaster recovery, and workload replay are similar to those for real-time migration.

Step 5 On the **Basic Information** page, click **Migration Logs** on the left.

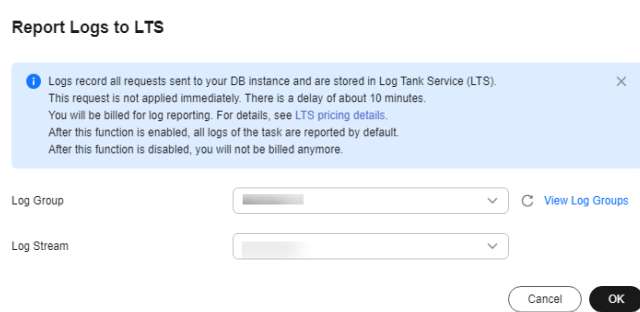
Step 6 Click  next to **Report Logs to LTS** in the upper part of the page.

Step 7 Select an LTS log group and log stream and click **OK**.

 **NOTE**

This request does not take effect immediately. There is a delay of about 10 minutes.


Figure 13-1 Enabling audit log reporting to LTS



Step 8 To disable or modify log reporting, click the toggle switch next to **Report Logs to LTS** or click **Edit** next to the **Report Logs to LTS** toggle switch.

- Modifying log reporting: Click **Edit** next to the **Report Logs to LTS** toggle switch. In the displayed dialog box, select the LTS log group and log stream again and click **OK**.
- Disabling log reporting: Click the toggle switch next to **Report Logs to LTS**. In the displayed dialog box, click **OK**.

Figure 13-2 Disabling log reporting to LTS



----End

13.2 Viewing and Downloading Logs

Scenarios

If you have enabled log reporting to LTS for a DRS task in [Log Reporting](#), you can analyze logs, search for logs, visualize logs, download logs, and view real-time logs on the LTS console.

Viewing Logs Reported to LTS


- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name. For details about LTS, see [Log Tank Service \(LTS\) User Guide](#).

Figure 13-3 Viewing log details




Table 13-1 Log field description

Name	Type	Description
_resource_id	String	Resource ID. The value is fixed to projectId for DRS.
_resource_name	String	Resource name. The value is fixed to DRS .
_service_type	String	Service type. The value is fixed to Data Replication Service .

----End

Downloading Logs Reported to LTS



- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name.
- Step 5** Click **Download** on the right to download logs. For details about LTS, see [Log Tank Service \(LTS\) User Guide](#).

Figure 13-4 Downloading logs



-----End

14 Operation Reference in Synchronization Scenarios

14.1 Kafka Message Format

Data synchronized to the Kafka cluster is stored in Avro, JSON, and JSON-C formats. For details about the data formats supported by different data flow scenarios, see [Table 14-1](#).

Table 14-1 Data formats

Data Flow	Avro	JSON	JSON-C
MySQL -> Kafka	Supported	Supported	Supported
Oracle -> Kafka	Supported	Supported	Not supported
DDS->Kafka	Not supported	Yes	Not supported
PostgreSQL -> Kafka	Supported	Supported	Not supported
GaussDB(for MySQL) -> Kafka	Supported	Supported	Supported
GaussDB primary/standby -> Kafka	Supported	Supported	Not supported
GaussDB distributed -> Kafka	Supported	Supported	Not supported
Microsoft SQL Server->Kafka	Supported	Supported	Not supported

Avro

For details about the schema definition in Avro format, see [record.rar](#). After data is synchronized to the Kafka cluster, parse data based on the definition of the Avro schema.

JSON

For details about the JSON format from MySQL and GaussDB (MySQL) to Kafka, see [Table 14-2](#). For details about the JSON format from DDS to Kafka, see [Table 14-3](#). For details about the JSON format from PostgreSQL, GaussDB, Microsoft SQL Server and Oracle to Kafka, see [Table 14-4](#).

Table 14-2 Parameters for synchronizing from MySQL to Kafka

Parameter	Description
mysqlType	Field name and type in the source table.
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
es	The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds.
ts	The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds.
database	Database name
table	Table name.
type	Operation type, for example, DELETE, UPDATE, INSERT, and DDL. For full synchronization, the value can be INIT or INIT_DDL.
isDdl	Whether the operation is a DDL operation.
sql	A DDL-defined SQL statement. The value is "".
sqlType	JDBC type of the fields in the source table.
data	The latest data, which is a JSON array. If the value of type is INSERT , the latest data is inserted. If the value of type is UPDATE , the latest data is updated.
old	Old data. If the value of type is UPDATE , the data is old. If the value of type is DELETE , the data is deleted.
pkNames	Primary key name

```
{
  "mysqlType":{
    "c11":"binary",
    "c10":"varchar",
    "c13":"text",
```

```

        "c12": "varbinary",
        "c14": "blob",
        "c1": "varchar",
        "c2": "varbinary",
        "c3": "int",
        "c4": "datetime",
        "c5": "timestamp",
        "c6": "char",
        "c7": "float",
        "c8": "double",
        "c9": "decimal",
        "id": "int"
    },
    "id": 27677,
    "es": 1624614713000,
    "ts": 1625058726990,
    "database": "test01",
    "table": "test ",
    "type": "UPDATE",
    "isDdl": false,
    "sql": "",
    "sqlType": {
        "c11": -2,
        "c10": 12,
        "c13": -1,
        "c12": -3,
        "c14": 2004,
        "c1": 12,
        "c2": -3,
        "c3": 4,
        "c4": 94,
        "c5": 93,
        "c6": 1,
        "c7": 6,
        "c8": 8,
        "c9": 3,
        "id": 4
    },
    "data": [
        {
            "c11": "[]",
            "c10": "Huawei Cloud huaweicloud",
            "c13": "asfiajhfi939-0239uoituqorjqirfojdjfrnioweioiwqjroqwjrowqjoiqgoiegnkjgoi23roiugou
            ofdug9u90weurtg103",
            "c12": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,
            103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,
            57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 101, 117, 114, 103, 57,
            101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114, 48, 52, 117, 48, 57, 53, 116, 117, 51, 48,
            57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",
            "c14": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,
            103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,
            57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 55, 57, 56, 52, 54, 53,
            52, 54, 54, 54, 49, 52, 54, 53, 33, 64, 35, 36, 37, 94, 42, 40, 41, 95, 41, 43, 95, 43, 124, 125, 34,
            63, 62, 58, 58, 101, 117, 114, 103, 57, 101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114,
            48, 52, 117, 48, 57, 53, 116, 117, 51, 48, 57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",
            "c1": "cf3f70a7-7565-44b0-ae3c-83bec549ea8e:104",
            "c2": "[]",
            "c3": "103",
            "c4": "2021-06-25 17:51:53",
            "c5": "1624614713.201",
            "c6": "!@#%$90weurtg103",
        }
    ]

```



```

        "c7":"10357.0",
        "c8":"1.2510357E7",
        "c9":"9874510357",
        "id":"104"
    }
],
"old":[
    {
        "c11": "",
        "c10": "Huawei Cloud huaweicloud",
        "c13": "asfiahfiah939-0239",
        "c12": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111, 103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48, 57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 101, 117, 114, 103, 57, 101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114, 48, 52, 117, 48, 57, 53, 116, 117, 51, 48, 57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",
        "c14": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111, 103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48, 57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 55, 57, 56, 52, 54, 53, 52, 54, 54, 54, 49, 52, 54, 53, 33, 64, 35, 36, 37, 94, 42, 40, 41, 95, 41, 43, 95, 43, 124, 125, 34, 63, 62, 58, 58, 101, 117, 114, 103, 57, 101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114, 48, 52, 117, 48, 57, 53, 116, 117, 51, 48, 57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",
        "c1": "cf3f70a7-7565-44b0-ae3c-83bec549ea8e:104",
        "c2": "",
        "c3": "103",
        "c4": "2021-06-25 17:51:53",
        "c5": "1624614713.201",
        "c6": "!@#%90weurtg103",
        "c7": "10357.0",
        "c8": "1.2510357E7",
        "c9": "9874510357",
        "id": "103"
    }
],
"pkNames":[
    "id"
]
}

```

Table 14-3 Parameters for synchronizing data from DDS to Kafka

Parameter	Description
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
op	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
dbType	Source database type: MongoDB
db	Database name.
coll	Collection name.
value	Change value of a record.
where	Change condition of a record.

Parameter	Description
recordType	Record type, such as insert, update, replace, and doc. update and replace indicate the UPDATE operation in op . doc indicates that the DELETE operation in op deletes document data instead of view data.
extra	Extended field. The value is the same as that of recordType and this parameter is used as an extended oplog record.
es	Commit time of a record. The value is a 13-digit Unix timestamp, in milliseconds.
ts	The time when the data is written to the destination Kafka. The value is a 13-digit Unix timestamp in milliseconds.
clusterTime	Timestamp of the oplog entry associated with the event. The value is in the format of timestamp:incr. timestamp is the Unix timestamp (unit: second), and incr is the command execution sequence in a second.

```
// insert operation
{
  "id": 256,
  "op": "INSERT",
  "dbType": "MongoDB",
  "db": "lxx",
  "coll": "lxx",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\"), \"c1\": \"baz\", \"tags\": [\"mongodb\", \"database\", \"NoSQL\"]}",
  "where": null,
  "recordType": "insert",
  "extra": "insert",
  "es": 1684315111439,
  "ts": 1684315111576,
  "clusterTime": "1684344064:1"
}

// replace operation
{
  "id": 340,
  "op": "UPDATE",
  "dbType": "MongoDB",
  "db": "lxx",
  "coll": "lxx",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\"), \"c1\": \"sss\"}",
  "where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "recordType": "replace",
  "extra": "replace",
  "es": 1684315951831,
  "ts": 1684315951961,
  "clusterTime": "1684344904:9"
}

// value update operation
{
  "id": 386,
  "op": "UPDATE",
```

```

"dbType": "MongoDB",
"db": "ljx",
"coll": "ljx",
"value": "{\"$set\": {\"c1\": \"aaa\"}}",
"where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
"recordType": "update",
"extra": "update",
"es": 1684316412008,
"ts": 1684316412146,
"clusterTime": "1684345365:1"
}

// key update operation
{
  "id": 414,
  "op": "UPDATE",
  "dbType": "MongoDB",
  "db": "ljx",
  "coll": "ljx",
  "value": "{\"$unset\": {\"c1\": true}, \"$set\": {\"column1\": \"aaa\"}}",
  "where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "recordType": "update",
  "extra": "update",
  "es": 1684316692054,
  "ts": 1684316692184,
  "clusterTime": "1684345648:1"
}

// remove operation
{
  "id": 471,
  "op": "DELETE",
  "dbType": "MongoDB",
  "db": "ljx",
  "coll": "ljx",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "where": null,
  "recordType": "doc",
  "extra": "doc",
  "es": 1684317252747,
  "ts": 1684317252869,
  "clusterTime": "1684346209:1"
}

```

Table 14-4 Parameters for synchronizing from other databases to Kafka

Parameter	Description
columnType	Field name and data type in the source table NOTE <ul style="list-style-type: none"> The data type does not contain the length and precision. This parameter is left blank when dbType is set to Oracle or Microsoft SQL Server.
dbType	Source database type
schema	Schema name.
opType	Operation type, such as DELETE, UPDATE, INSERT, and DDL.

Parameter	Description
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
es	The source DB engine types are as follows: GaussDB Primary/Standby: commit time of the current transaction. The value is a 13-digit Unix timestamp, in milliseconds. GaussDB Distributed: commit time of the current transaction. The value is a 13-digit Unix timestamp, in milliseconds. PostgreSQL: commit time of the previous transaction. The value is a 13-digit Unix timestamp, in milliseconds. Oracle: commit time of a record. The value is a 13-digit Unix timestamp, in milliseconds. Microsoft SQL Server: commit time of a record. The value is a 13-digit Unix timestamp, in milliseconds.
ts	The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds.
database	Database name. This parameter is left blank when dbType is set to Oracle .
table	Table name.
type	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
isDdl	Whether the operation is a DDL operation.
sql	A DDL-defined SQL statement. The value is "".
sqlType	JDBC type of the fields in the source table.
data	The latest data, which is a JSON array. If the value of type is INSERT , the latest data is inserted. If the value of type is UPDATE , the latest data is updated.
old	Old data. If the value of type is UPDATE , the data is old. If the value of type is DELETE , the data is deleted.
pkNames	Primary key name

```
{
  "columnType": {
    "timestamp_column": "timestamp without time zone",
    "tstzrange_column": "tstzrange",
    "int4range_column": "int4range",
    "char_column": "character",
    "jsonb_column": "json",
    "boolean_column": "boolean",
    "bit_column": "bit",
    "smallint_column": "smallint",
    "bytea_column": "bytea"
  },
}
```

```
"dbType": "GaussDB Primary/Standby",
"schema": "schema01",
"opType": "UPDATE",
"id": 332,
"es": 1639626187000,
"ts": 1639629261915,
"database": "database01",
"table": "table01",
"type": "UPDATE",
"isDdl": false,
"sql": "",
"sqlType": {
  "timestamp_column": 16,
  "tstzrange_column": 46,
  "int4range_column": 42,
  "char_column": 9,
  "jsonb_column": 22,
  "boolean_column": 8,
  "bit_column": 20,
  "smallint_column": 2,
  "bytea_column": 15
},
"data": [
  {
    "timestamp_column": "2021-12-16 12:31:49.344365",
    "tstzrange_column": "(\"2010-01-01 14:30:00+08\", \"2010-01-01 15:30:00+08\")",
    "int4range_column": "[11,20)",
    "char_column": "g",
    "jsonb_column": "{\"key1\": \"value1\", \"key2\": \"value2\"}",
    "boolean_column": "false",
    "bit_column": "1",
    "smallint_column": "12",
    "bytea_column": "62797465615f64617461"
  }
],
"old": [
  {
    "timestamp_column": "2014-07-02 06:14:00.742",
    "tstzrange_column": "(\"2010-01-01 14:30:00+08\", \"2010-01-01 15:30:00+08\")",
    "int4range_column": "[11,20)",
    "char_column": "g",
    "jsonb_column": "{\"key1\": \"value1\", \"key2\": \"value2\"}",
    "boolean_column": "true",
    "bit_column": "1",
    "smallint_column": "12",
    "bytea_column": "62797465615f64617461"
  }
],
"pkNames": null
}
```

JSON-C

JSON-C is similar to JSON. The difference lies in the delete operation. JSON data is stored in old, and JSON-C is stored in data. Data of the timestamp type is converted into a character string in the format of yyyy-mm-dd hh:mm:ss.

For details, see [Table 14-5](#).

Table 14-5 JSON-C parameter description

Parameter	Description
mysqlType	Field name and type in the source table.
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
es	The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds.
ts	The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds.
database	Database name. For the Oracle database, set this parameter to schema .
table	Table name.
type	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
isDdl	Whether the operation is a DDL operation.
sql	A DDL-defined SQL statement. The value is "".
sqlType	JDBC type of the fields in the source table.
data	Latest data, which is a JSON array. If type is set to INSERT , this parameter indicates the latest inserted data. If type is set to UPDATE , this parameter indicates the latest updated data. If type is set to DELETE , this parameter indicates the deleted data.
old	Old data. If type is set to UPDATE , the value indicates the data before update. If type is set to INSERT , the value is null .
pkNames	Primary key name

Common Escape Characters in JSON

Table 14-6 Escape Character

Character	Escape character
<	\u003c
=	\u003d
>	\u003e
&	\u0026amp;
'	\u0027

14.2 Kafka Authentication

PLAINTEXT

No security authentication mode is available. You only need to enter the IP address and port for connection.

Figure 14-1 PLAINTEXT

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

SASL_PLAINTEXT

The SASL mechanism is used to connect to Kafka, and you need to configure SASL parameters.

Figure 14-2 SASL_PLAINTEXT

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

SASL Mechanisms ⓘ

Username

Password ⓘ

Table 14-7 Parameter settings

Parameter	Description
SASL Mechanisms	<p>SASL is used by client. The following four items are supported. Kafka server uses the GSSAPI mechanism by default. For details, see the SASL mechanisms.</p> <ul style="list-style-type: none"> ● GSSAPI ● PLAIN ● SCRAM-SHA-256 ● SCRAM-SHA-512

Parameter	Description
Token Delegation	Whether an agency token is used for authentication. This option is available when SCRAM-SHA-256 or SCRAM-SHA-512 is selected for SASL Mechanisms .
Username	Username for logging in to the database
Password	Password for the username

SSL

SSL is used to encrypt the connection to Kafka. Related parameters need to be configured.

Figure 14-3 SSL

Destination Database

IP Address ⓘ

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

Truststore Certificate Select

Truststore Certificate Password ⓘ

Endpoint Identification Algorithm ⓘ

Mutual SSL Authentication

Keystore Certificate Select

Keystore Certificate Password ⓘ

Keystore Private Key Password ⓘ

Table 14-8 Parameter settings

Parameter	Description
Truststore Certificate	SSL certificate with the file name extension .jks. For details about how to generate a certificate file, see SSL Certificate .
Truststore Certificate Password	Password of the certificate

Parameter	Description
Endpoint Identification Algorithm	Endpoint identification algorithm for verifying the host name of the server using the server certificate. This parameter is optional. If this parameter is left blank, host name verification is disabled.
Mutual SSL Authentication	Mutual SSL Authentication
Keystore Certificate	If mutual SSL authentication is enabled, you need to upload the mutual SSL authentication certificate with the file name extension .jks.
Keystore Certificate Password	Password of the mutual SSL authentication certificate. This option is available if mutual SSL authentication is enabled.
Keystore Private Key Password	(Optional) Password of the private key in the keystore certificate.

SASL_SSL

If the SASL and SSL are used, configure SSL and SASL parameters. For details, see [SASL_PLAINTEXT](#) and [SSL](#).

Figure 14-4 SASL_SSL

Destination Database

IP Address ⓘ
Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

Security Protocol ⓘ

SASL Mechanisms ⓘ

Username

Password

Truststore Certificate

Truststore Certificate Password ⓘ

Endpoint Identification Algorithm ⓘ

Mutual SSL Authentication

Keystore Certificate

Keystore Certificate Password ⓘ

Keystore Private Key Password ⓘ

This button is available only after the replication instance is created successfully.

14.3 Synchronizing Sequence Values for a Synchronization Task with GaussDB Distributed Serving as the Source Database

A task with GaussDB Distributed serving as the source database does not support sequence value synchronization. If the tables to be synchronized contain associated sequences, you need to manually synchronize the sequence values after the task is complete.

Procedure

Step 1 Log in to the distributed GaussDB instance as the user used when you tested connectivity between the DRS instance and the distributed GaussDB instance.

Step 2 Run the following statement to query the tables that have management sequences and their associated sequences in the database:

```
set search_path to '';select d.refobjid::regclass::text as tablename, d.objid::regclass::text as seqname from pg_depend d where d.refclassid='pg_class'::regclass and d.objid in (select oid from pg_class where relkind in ('S','L')) UNION select a.adrelid::regclass::text,d.refobjid::regclass::text from pg_attrdef a JOIN pg_depend d ON (a.oid=d.objid) where d.refobjid in (select oid from pg_class where relkind in ('S','L')) and d.classid='pg_attrdef'::regclass order by tablename, seqname;
```

Step 3 For the sequence associated with each table to be synchronized, run the following statement on the CN node of the source database as a user with the MONADMIN or SYSADMIN permission to query **next_new_val** corresponding to all nodes:

```
execute direct on all $$select last_value, increment_by, cache_value, (last_value + increment_by * cache_value) as next_new_val from '<seqname>'$$;
```

In the preceding command, **<seqname>** indicates the name of the sequence with schema in the query result of [Step 2](#). The value of **next_new_val** in the query result is **last_value + increment_by * cache_value**. Ensure that the sequence value is unique.

Step 4 For the sequence associated with each table to be synchronized, run the following statement to set a new sequence value in the destination database:

```
select setval('<seqname>', <target_value>);
```

In the preceding command, **<target_value>** indicates the maximum value (auto-increment sequence, **increment_by** > 0) or minimum value (auto-decrement, **increment_by** < 0) of **next_new_val** queried in [Step 3](#).

Step 5 Check the sequence value of the destination database.

For the sequence associated with each table to be synchronized, run the following statement to obtain the new value of the sequence in the destination database:

```
select nextval('<seqname>');
```

Check whether the new value is greater than or equal to (auto-increment sequence, **increment_by** > 0) or less than or equal to (auto-decrement, **increment_by** < 0) the sequence values of all nodes in the source database.

----End

14.4 Forcibly Stopping Synchronization of GaussDB Distributed

This section describes how to clear the streaming replication slots of the source GaussDB distributed database after the incremental or full+incremental synchronization task is forcibly stopped.

The naming rule of a replication slot depends on the database kernel version (you can run the **select working_version_num();** command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Prerequisites

Common users do not have the permission to perform the execute direct operation. To delete streaming replication slots, contact GaussDB O&M personnel.

Procedure

Step 1 Log in to each primary DN node of the GaussDB distributed instance as the user used when you tested the connectivity between the DRS instance and the GaussDB distributed instance.

Step 2 Run the following statement to query the streaming replication slot name of the database object selected for the synchronization task:

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

Step 3 Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

Step 4 Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where database = 'database';
```

If the query result is empty, the streaming replication slot is deleted.

- Step 5** Repeat the preceding operations to ensure that the streaming replication slot on each primary DN is deleted.

----End

14.5 Forcibly Stopping Synchronization of GaussDB Primary/Standby

This section describes how to clear the streaming replication slots of the source GaussDB primary/standby database after the synchronization task is forcibly stopped.

The naming rule of a replication slot depends on the database kernel version (you can run the **select working_version_num();** command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is `replication_slot_unique_ID`. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is `drs_unique_ID`. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Prerequisites

Common users do not have the permission to perform the execute direct operation. To delete streaming replication slots, contact GaussDB O&M personnel.

Procedure

- Step 1** Log in to the GaussDB primary/standby instance as the user used when you tested connectivity between the DRS instance and the GaussDB primary/standby instance.

- Step 2** Run the following statement to query the streaming replication slot name of the database object selected for the synchronization task:

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

- Step 3** Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

- Step 4** Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where database = 'database';
```

If the query result is empty, the streaming replication slot is deleted.

----End

14.6 Forcibly Stopping Synchronization of PostgreSQL

This section describes how to clear the logical replication slot of the source database, how to synchronize sequence values, and how to reset the sequence values in the destination database when the source database cannot be connected after the PostgreSQL synchronization task is forcibly stopped.

The naming rule of a replication slot is `drs_unique_ID`. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Clearing the Logical Replication Slot of the Source Database

- Step 1** Log in to the source database as the source database user used in the synchronization task.

- Step 2** Query the name of the streaming replication slot of the database object selected in the synchronization task.

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

- Step 3** Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

- Step 4** Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where slot_name = 'slot_name';
```

If the query result is empty, the streaming replication slot is deleted.

----End

Synchronizing Sequence Values

If sequence objects are not synchronized or the destination database is GaussDB, skip this section.

- Step 1** Use a high-privilege account (with the USAGE permission for all sequences) to connect to the source database and run the following statement:

```
select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(n.nspname))||','||quote_ident(c.relname))||', '||nextval(c.oid)||');' as sqls from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' and not (c.relname='hwdrs_ddl_info_id_seq' and n.nspname='public') order by n.nspname, c.relname;
```

The query result is the SQL statement that needs to be executed in the destination database.

- Step 2** Log in to the destination database as the destination database user used in the synchronization task and run the SQL statement queried in **Step 1** in the destination database.

- Step 3** Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' order by 1,2;
```

----End

Resetting Sequence Values in the Destination Database

If the source database failed and cannot be connected, you can reset the sequence values related to automatic increment or decrement columns in the destination database. If the source database can be connected, skip this section.

- Step 1** Log in to the destination database as the destination database user used in the synchronization task.

- Step 2** Run the following statement to query the SQL statement for resetting the sequence value corresponding to the sequence that uses nextval as the default value of the table column:

```
set search_path to ''; select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(s.sequence_schema))||','||quote_ident(s.sequence_name))||', (SELECT '||case when s.increment::int<0 then 'min(' else 'max(' end||quote_ident(c.column_name)||')' ||case when s.increment::int<0 then '-1' else '+1' end||' FROM '||quote_ident(c.table_schema)||','||quote_ident(c.table_name)||');' as sqls from information_schema.columns c join information_schema.sequences s on (position(quote_literal(quote_ident(s.sequence_schema))||','||quote_ident(s.sequence_name))||':regclass' in c.column_default) > 0) where c.data_type in ('bigint', 'int', 'integer', 'smallint', 'numeric', 'real', 'double precision', 'double') and c.column_default like 'nextval(%%%' order by s.sequence_schema, s.sequence_name;
```

The query result is the SQL statement that needs to be executed in the destination database.

- Step 3** If the source database version is earlier than 10.0, skip this step. If the source database version is 10.0 or later, run the following statement in the destination database to query the SQL statement for resetting the sequence value corresponding to the additional column of the table identity column:

```
set search_path to ''; select 'SELECT pg_catalog.setval('||quote_literal(seqname))||', (SELECT '||case when increment::int<0 then 'min(' else 'max(' end||colname||')' ||case when increment::int<0 then '-1' else '+1' end||' FROM '||tablename||');' as sqls from (select objid::regclass::text, refobjid::regclass::text, (pg_identify_object(refclassid,refobjid,refobjsubid)).identity, (pg_sequence_parameters(objid)).increment from pg_depend where deptype='i' and refobjsubid>0 and objid in (select c.oid from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind='S' and n.nspname !~ '^pg_' and n.nspname <> 'information_schema')) p(seqname,tablename,colname,increment);
```

The query result is the SQL statement that needs to be executed in the destination database.

Step 4 Run the SQL statements queried in **Step 2** and **Step 3** in the destination database.

Step 5 Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid  
where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname<>'information_schema' order by 1,2;
```

----End

14.7 Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL

This section describes how to perform real-time synchronization from PostgreSQL to RDS for PostgreSQL. You can create event triggers and functions in the source database to obtain the DDL information of the source database, and then synchronize DDL operations to the destination database during the incremental synchronization phase.

Prerequisites

- The following DDL operations are supported:
 - Table-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), COMMENT ON COLUMN, and COMMENT ON TABLE
 - Database-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA/TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW

CAUTION

- Table-level synchronization: If data is inserted into a renamed table, the data will not be synchronized to the destination database.
 - Database-level synchronization: Tables that are created not using the CREATE TABLE statement in the source database will not be synchronized to the destination database. For example, you run CREATE TABLE AS to create a table or call a function to create a table.
 - DDL statements starting with comments cannot be synchronized and are ignored.
 - DDL statements executed in functions and stored procedures cannot be synchronized and are ignored.
-
- If the source and destination databases are of different versions, use SQL statements that are compatible with both the source and destination databases to perform DDL operations. For example, if the source database is PostgreSQL 11 and the destination database is PostgreSQL 12, run the following statement to change the column type from char to int:

```
alter table tablename alter column columnname type int USING columnname::int;
```
 - Check whether a table named **hwdrs_ddl_info**, a function named **hwdrs_ddl_function()**, and a trigger named **hwdrs_ddl_event** exist in the source database in public mode. If they exist, delete them.
 - During database-level synchronization, if a table without a primary key is created, run the following command to set the replication attribute of the table without a primary key to full.

```
alter table tablename replica identity full;
```

Procedure

- If the source database is a self-managed PostgreSQL database or a database on another cloud platform, perform the following steps:
 - a. Connect to the database to be synchronized as a user who has permission to create event triggers.
 - b. Run the following statements to create a table for storing DDL information:

```
DROP TABLE IF EXISTS public.hwdrs_ddl_info;
DROP SEQUENCE IF EXISTS public.hwdrs_ddl_info_id_seq;
CREATE TABLE public.hwdrs_ddl_info(
  id          bigserial primary key,
  ddl        text,
  username   varchar(64) default current_user,
  txid       varchar(16) default txid_current()::varchar(16),
  tag        varchar(64),
  database   varchar(64) default current_database(),
  schema     varchar(64) default current_schema,
  client_addr varchar(64) default inet_client_addr(),
  client_port integer default inet_client_port(),
  event_time timestamp default current_timestamp
);
```
 - c. Run the following statements to create a function:

```
CREATE OR REPLACE FUNCTION public.hwdrs_ddl_function()
  RETURNS event_trigger
  LANGUAGE plpgsql
  SECURITY INVOKER
AS $$
```



```
declare ddl text;
declare real_num int;
declare max_num int := 50000;
begin
  if (tg_tag in ('CREATE TABLE','ALTER TABLE','DROP TABLE','CREATE SCHEMA','CREATE SEQUENCE','ALTER SEQUENCE','DROP SEQUENCE','CREATE VIEW','ALTER VIEW','DROP VIEW','CREATE INDEX','ALTER INDEX','DROP INDEX','COMMENT')) then
    select current_query() into ddl;
    insert into public.hwdrs_ddl_info(ddl, username, txid, tag, database, schema, client_address, client_port, event_time)
    values (ddl, current_user, cast(txid_current() as varchar(16)), tg_tag, current_database(), current_schema, inet_client_addr(), inet_client_port(), current_timestamp);
    select count(id) into real_num from public.hwdrs_ddl_info;
    if real_num > max_num then
      if current_setting('server_version_num')::int < 100000 then
        delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and not exists (select 0 from pg_locks l join pg_database d on l.database=d.oid where d.datname=current_catalog and pid <> pg_backend_pid() and locktype='relation' and relation=to_regclass('public.hwdrs_ddl_info_pkey')::oid and mode='RowExclusiveLock');
      else
        delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and (xmax=0 or coalesce(txid_status(xmax::text)::bigint, '') <> 'in progress');
      end if;
    end if;
  end if;
end;
$$;
```

- d. Run the following statements to grant necessary permissions to the objects created in **b** and **c**:

```
GRANT USAGE ON SCHEMA public TO public;
GRANT SELECT,INSERT,DELETE ON TABLE public.hwdrs_ddl_info TO public;
GRANT SELECT,USAGE ON SEQUENCE public.hwdrs_ddl_info_id_seq TO public;
GRANT EXECUTE ON FUNCTION public.hwdrs_ddl_function() TO public;
```

- e. Run the following statement to create a DDL event trigger:

```
CREATE EVENT TRIGGER hwdrs_ddl_event ON ddl_command_end EXECUTE PROCEDURE public.hwdrs_ddl_function();
```

- f. Run the following statement to set the created event trigger to enable:

```
ALTER EVENT TRIGGER hwdrs_ddl_event ENABLE ALWAYS;
```

- g. Return to the DRS console and create a PostgreSQL to RDS PostgreSQL synchronization task.

- h. After the synchronization task is complete, run the following statements to delete the created tables, functions, and triggers.

```
DROP EVENT TRIGGER hwdrs_ddl_event;
DROP FUNCTION public.hwdrs_ddl_function();
DROP TABLE public.hwdrs_ddl_info;
```

- If the source database is RDS for PostgreSQL, perform the following steps:

- a. Run the following statements to delete the created objects:

```
DROP EVENT TRIGGER IF EXISTS hwdrs_ddl_event;
DROP FUNCTION IF EXISTS public.hwdrs_ddl_function();
DROP TABLE IF EXISTS public.hwdrs_ddl_info;
```

- b. Run the following statement as the **root** user to create a DDL plug-in:

```
SELECT CONTROL_EXTENSION('create', 'rds_hwdrs_ddl');
```

- c. Run the following statements to update the function:

```
CREATE OR REPLACE FUNCTION public.hwdrs_ddl_function()
  RETURNS event_trigger
  LANGUAGE plpgsql
  SECURITY INVOKER
AS $BODY$
declare ddl text;
declare real_num int;
```

```
declare max_num int := 50000;
begin
  if (tg_tag in ('CREATE TABLE','ALTER TABLE','DROP TABLE','CREATE SCHEMA','CREATE
SEQUENCE','ALTER SEQUENCE','DROP SEQUENCE','CREATE VIEW','ALTER VIEW','DROP
VIEW','CREATE INDEX','ALTER INDEX','DROP INDEX','COMMENT')) then
    select current_query() into ddl;
    insert into public.hwdrs_ddl_info(ddl, username, txid, tag, database, schema, client_address,
client_port, event_time)
    values (ddl, current_user, cast(txid_current() as varchar(16)), tg_tag, current_database(),
current_schema, inet_client_addr(), inet_client_port(), current_timestamp);
    select count(id) into real_num from public.hwdrs_ddl_info;
    if real_num > max_num then
      if current_setting('server_version_num')::int<100000 then
        delete from public.hwdrs_ddl_info where id<(select min(id)+1000 from
public.hwdrs_ddl_info) and not exists (select 0 from pg_locks l join pg_database d on
l.database=d.oid where d.datname=current_catalog and pid<>pg_backend_pid() and
locktype='relation' and relation=to_regclass('public.hwdrs_ddl_info_pkey')::oid and
mode='RowExclusiveLock');
      else
        delete from public.hwdrs_ddl_info where id<(select min(id)+1000 from
public.hwdrs_ddl_info) and (xmax=0 or coalesce(txid_status(xmax::text::bigint), '')<>'in
progress');
      end if;
    end if;
  end if;
end;
$BODY$;
```

- d. Return to the DRS console and create a synchronization task from PostgreSQL to RDS for PostgreSQL.
- e. After the synchronization task is complete, run the following statement to delete the created tables, functions, and triggers.

```
SELECT CONTROL_EXTENSION('drop', 'rds_hwdrs_ddl');
```

15 Appendix

15.1 General Usage Suggestions

Before using DRS, you need to understand the following:

- The destination database version must be the same as or later than the source database version.
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to be synchronized successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
- Data-Level Comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.
- If many-to-one synchronization is involved, you need to understand the following:
 - If you create a many-to-one synchronization task, refer to [Constraints and Operation Suggestions on Many-to-One Scenario](#).
 - For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail.
- After data synchronization is complete, you need to understand the following:
 - To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

15.2 DRS Pre-Check Items

A DRS synchronization process consists of four phases: task startup, full migration, incremental migration, and task completion. A single full migration task contains three phases. To ensure smooth synchronization, DRS automatically scans and checks parameters and objects in the source and destination databases before starting a synchronization task.

Overview

View pre-check items based on the source database type of your synchronization task.

- The following lists the pre-check items for synchronization tasks with MySQL serving as the source.
 - [MySQL->MySQL](#)
 - [MySQL->PostgreSQL](#)
 - [MySQL -> GaussDB Distributed](#)
 - [MySQL -> GaussDB Primary/Standby](#)
 - [MySQL->GaussDB\(DWS\)](#)

- **MySQL->GaussDB(for MySQL)**
- **MySQL->MariaDB**
- The following lists the pre-check items for synchronization tasks with PostgreSQL serving as the source.
 - **PostgreSQL->PostgreSQL**
 - **PostgreSQL->GaussDB(DWS)**
 - **PostgreSQL -> GaussDB Primary/Standby**
 - **PostgreSQL -> GaussDB Distributed**
- The following lists the pre-check items for synchronization tasks with Oracle serving as the source.
 - **Oracle->MySQL**
 - **Oracle->GaussDB(for MySQL)**
 - **Oracle -> GaussDB Primary/Standby**
 - **Oracle -> GaussDB Distributed**
 - **Oracle->GaussDB(DWS)**
 - **Oracle->DDM**
 - **Oracle->PostgreSQL**
- The following lists the pre-check items for synchronization tasks with DDM serving as the source.
 - **DDM->MySQL**
 - **DDM->GaussDB(DWS)**
 - **DDM->DDM**
- The following lists the pre-check items for synchronization tasks with DB2 for LUW serving as the source.
 - **DB2 for LUW -> GaussDB Primary/Standby**
 - **DB2 for LUW -> GaussDB Distributed**
 - **DB2 for LUW->GaussDB(DWS)**
- The following lists the pre-check items for synchronization tasks with TiDB serving as the source.
 - **TiDB->GaussDB(for MySQL)**
- The following lists the pre-check items for synchronization tasks with Microsoft SQL Server serving as the source.
 - **Microsoft SQL Server->GaussDB(DWS)**
 - **Microsoft SQL Server -> GaussDB Primary/Standby**
 - **Microsoft SQL Server -> GaussDB Distributed**
 - **Microsoft SQL Server->Microsoft SQL Server**
- The following lists the pre-check items for synchronization tasks with MongoDB serving as the source.
 - **MongoDB->DDS**
- The following lists the pre-check items for synchronization tasks with MariaDB serving as the source.
 - **MariaDB->MySQL**

- [MariaDB->GaussDB\(for MySQL\)](#)
- The following lists the pre-check items for synchronization tasks with GaussDB(for MySQL) serving as the source.
[GaussDB\(for MySQL\)->GaussDB\(for MySQL\)](#)

MySQL->MySQL

Table 15-1 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> • Full synchronization requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT • Full+incremental synchronization and incremental synchronization require the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> • The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES <p>The root account of the RDS for MySQL DB instance has the preceding permissions by default.</p> <ul style="list-style-type: none"> • If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. 	For details, see Which MySQL Permissions Are Required for DRS?
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.
	Destination database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database.
	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured.
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL.
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the source database is too small, the synchronization task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small.
If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the destination database is too small, data cannot be written to the destination database and the full synchronization fails.		For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small.	

Category	Check Item	Check Item Details	Solution to Failure
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental synchronization	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database .
	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	<p>For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.</p>
	Table fields	<p>If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.</p>	-
	Additional columns	<p>In an incremental synchronization task, if no additional column is added to the destination database, the task fails.</p>	<p>For details, see Checking the Additional Column of the Destination Database.</p>
Destination database	Storage space	<p>The destination DB instance must have sufficient storage space.</p>	<p>For details, see Checking Whether the Destination Database Has Sufficient Storage Space.</p>
	Status	<p>The destination DB instance is running properly.</p>	-
Consistency	Character set	<p>The character set of the destination database must be the same as that of the source database.</p>	<p>For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.</p>

Category	Check Item	Check Item Details	Solution to Failure
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.

Category	Check Item	Check Item Details	Solution to Failure
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Database parameters and the system database cannot be synchronized. 	-
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects.

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration.
	Foreign key references	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.
	Duplicate names	If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<>\'\" • The source table and view names cannot contain non-ASCII characters, or the following characters: '<>\'\" • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured .
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

MySQL->PostgreSQL

Table 15-2 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following minimum permission: SELECT Full+incremental synchronization requires the following minimum permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	The account of the RDS for PostgreSQL instance has the permissions by default.	-
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
	Destination database versions	Versions 9.5, 9.6, 10, and 11 are supported.	For details, see Supported Databases .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the source database is too small, the synchronization task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small .

Category	Check Item	Check Item Details	Solution to Failure
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled.
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based.
	server_id value	During an incremental synchronization, the server_id value of the MySQL source database must be set. <ul style="list-style-type: none"> If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.
	Table fields	<ul style="list-style-type: none"> If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. The following table field types are not supported in the source database: xml, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and json. 	-
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.

Category	Check Item	Check Item Details	Solution to Failure
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Table structures	The table structure of the destination database must be the same as that of the source database.	For details, see Checking Whether Tables Structures Are Consistent.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Only table structures, table data, and indexes can be synchronized. ● Only MyISAM and InnoDB tables can be synchronized. ● Other database objects such as stored procedures cannot be synchronized. 	-
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Primary keys	Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario.	-

Category	Check Item	Check Item Details	Solution to Failure
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database .
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same .
	Database table naming rules	The names of the source databases and tables cannot contain non-ASCII characters, or the following characters: '<>/'	-
	Object name length	The destination database object name can contain a maximum of 63 characters.	For details, see Checking Whether the Length of the Source Database Object Names Exceeds the Limit .
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured .
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .

Category	Check Item	Check Item Details	Solution to Failure
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

MySQL -> GaussDB Primary/Standby

Table 15-3 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the object, and grant the CREATE and USAGE permissions of the schema to the user. Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the object, and grant the SELECT, UPDATE, INSERT, ALTER, INDEX, and DELETE permissions for all tables in the SCHEMA to the user. 	-
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the <code>max_allowed_packet</code> value of the source database is too small, the synchronization task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .

Category	Check Item	Check Item Details	Solution to Failure
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .
	Table fields	<p>The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon.</p>	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Database compatibility type	<p>The compatibility mode of the destination database must be consistent with that of the source database.</p>	For details, see Checking the Database Compatibility Type .

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Table structures	The table structure of the destination database must be the same as that of the source database.	For details, see Checking Whether Tables Structures Are Consistent.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Only table structures, table data, and indexes can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Other database objects such as stored procedures cannot be synchronized. • MySQL tables containing virtual columns cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.
	Database table naming rules	<ul style="list-style-type: none"> The names of the source databases and tables cannot contain non-ASCII characters, or the following characters: .<'>/\ Object names will be converted to lowercase letters after being synchronized to the destination database. Ensure that the selected source database tables do not contain tables with the same name but different letter cases. 	-
	Object name length	The destination database object name can contain a maximum of 63 characters.	For details, see Checking Whether the Length of the Source Database Object Names Exceeds the Limit.
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects.
	Destination database	The mapped database configured for the task must exist in the destination database.	-

Category	Check Item	Check Item Details	Solution to Failure
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MySQL->GaussDB(DWS)

Table 15-4 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE.	-
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.
	Destination database versions	Versions 8.1.3 and 8.2.0 are supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the source database is too small, the synchronization task may fail.	For details, see Checking Whether the max_allowed_packet Value of the Source Database Is Too Small .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .

Category	Check Item	Check Item Details	Solution to Failure
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .
	Table fields	<p>The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon.</p>	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .

Category	Check Item	Check Item Details	Solution to Failure
		If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-
	Database compatibility type	The compatibility mode of the destination database must be consistent with that of the source database.	For details, see Checking the Database Compatibility Type .
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space .
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent .
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same .
	Table structures	The table structure of the destination database must be the same as that of the source database.	For details, see Checking Whether Tables Structures Are Consistent .

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, indexes, and constraints (primary key, null, and non-null) can be synchronized. • MyISAM and InnoDB tables can be synchronized. • Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, and unique indexes cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. 	-
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> The names of the source databases and tables cannot contain non-ASCII characters, or the following characters: .<'>/\ During database-level synchronization, tables with the same name but different letter cases cannot be created in the source database during incremental synchronization. Object names will be converted to lowercase letters after being synchronized to the destination database. Ensure that the selected source database tables do not contain tables with the same name but different letter cases. 	-
	Object name length	The destination database object name can contain a maximum of 63 characters.	For details, see Checking Whether the Length of the Source Database Object Names Exceeds the Limit.
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects.
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.

Category	Check Item	Check Item Details	Solution to Failure
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MySQL->GaussDB(for MySQL)

Table 15-5 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	The user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES The root account of the GaussDB(for MySQL) DB instance has the preceding permissions by default. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. 	-
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.
	Destination database versions	Version 8.0 is supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database.
	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured.
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL.
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the source database is too small, the synchronization task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small.
		If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the destination database is too small, data cannot be written to the destination database and the full synchronization fails.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small.

Category	Check Item	Check Item Details	Solution to Failure
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental synchronization	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database .
	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	<p>For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.</p>
	Table fields	<p>If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.</p>	-
	Additional columns	<p>Check whether you need to add the missing additional columns to the destination database tables.</p>	<p>For details, see Checking the Additional Column of the Destination Database.</p>
Destination database	Status	<p>The destination DB instance is running properly.</p>	-
Consistency	Character set	<p>The character set of the destination database must be the same as that of the source database.</p>	<p>For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.</p>

Category	Check Item	Check Item Details	Solution to Failure
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.

Category	Check Item	Check Item Details	Solution to Failure
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Database parameters and the system database cannot be synchronized. 	-
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects.

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration .
	Foreign key references	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database .
	Duplicate names	If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same .
	Database table naming rules	<ul style="list-style-type: none"> The names of the source databases, tables, and views cannot contain non-ASCII characters, or the following characters: '<'>/'\ The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-

Category	Check Item	Check Item Details	Solution to Failure
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MySQL->MariaDB

Table 15-6 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT Full+incremental synchronization and incremental synchronization require the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX <p>The root account of the RDS for MariaDB instance has the preceding permissions by default.</p>	-

Category	Check Item	Check Item Details	Solution to Failure
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
	Destination database versions	Version 10.5 is supported.	For details, see Supported Databases .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .	
Incremental synchronization	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database .
	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .

Category	Check Item	Check Item Details	Solution to Failure
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based.
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database.
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.
	Table fields	If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-

Category	Check Item	Check Item Details	Solution to Failure
	Additional columns	In an incremental synchronization task, if no additional column is added to the destination database, the task fails.	For details, see Checking the Additional Column of the Destination Database.
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.

Category	Check Item	Check Item Details	Solution to Failure
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Database parameters and the system database cannot be synchronized. 	-
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects.

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration.
	Foreign key references	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.
	Duplicate names	If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<>/' • The source table and view names cannot contain non-ASCII characters, or the following characters: '<>/' • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MySQL -> GaussDB Distributed

Table 15-7 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> • Full synchronization: SELECT. • Full+incremental synchronization and incremental synchronization: SELECT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> • Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. • Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. • Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the max_allowed_packet value of the source database is too small, the synchronization task may fail.	For details, see Checking Whether the max_allowed_packet Value of the Source Database Is Too Small .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .

Category	Check Item	Check Item Details	Solution to Failure
	Binlog retention period	<p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental synchronization, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .
	Table fields	<p>The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and interval.</p>	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Database compatibility type	<p>The compatibility mode of the destination database must be consistent with that of the source database.</p>	For details, see Checking the Database Compatibility Type .

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Table structures	The table structure of the destination database must be the same as that of the source database.	For details, see Checking Whether Tables Structures Are Consistent.

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Only table structures, table data, and indexes can be synchronized. ● Only tables with primary keys can be synchronized. Tables without primary keys cannot be synchronized. ● Only MyISAM and InnoDB tables can be synchronized. ● Other database objects such as stored procedures cannot be synchronized. ● Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. 	-
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> The names of the source databases and tables cannot contain non-ASCII characters, or the following characters: .<'>\/ Object names will be converted to lowercase letters after being synchronized to the destination database. Ensure that the selected source database tables do not contain tables with the same name but different letter cases. 	-
	Object name length	The destination database object name can contain a maximum of 63 characters.	For details, see Checking Whether the Length of the Source Database Object Names Exceeds the Limit.
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects.
	Destination database	The mapped database configured for the task must exist in the destination database.	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

PostgreSQL->PostgreSQL

Table 15-8 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following permissions: The CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, SELECT permission for sequences, and SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords) Full+incremental synchronization and incremental synchronization require the following permissions: The CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, SELECT permission for sequences, SELECT permission for system table pg_catalog.pg_authid (used for synchronizing user passwords), UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create a replication connection 	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> ● Database-level synchronization: <ul style="list-style-type: none"> – If the destination database is not PostgreSQL, the CREATEDB permission is required. – If the destination database is PostgreSQL, the CONNECT and CREATE permissions on PostgreSQL databases and the USAGE and CREATE permissions on public schemas are required. ● Table-level synchronization: <ul style="list-style-type: none"> – To synchronize databases, the CREATEDB permission is required. – To synchronize schemas, the CONNECT and CREATE permissions for the database that contains the schemas are required. – To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the objects are required. ● User synchronization: The CREATEROLE permission is required. ● User permissions synchronization: The default privilege cannot be modified. Otherwise, the object permissions of the destination database may be inconsistent with those of the source database. 	-
Versions	Source database versions	Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .
	Destination database versions	Versions 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.
Parameters	Performance parameters	The BLOCK_SIZE value of the destination database must be greater than or equal to that of the source database.	For details, see Checking Whether the BLOCK_SIZE Value of the Source Database Is the Same as That of the Destination Database.
		The PASSWORD_ENCRYPTION values of the source and destination databases must be the same.	-
		The max_replication_slots value of the source database must be greater than the number of used replication slots.	For details, see Checking Whether the MAX_REPLICATION_SLOTS Value in the Source Database Is Correct.
		The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value.	For details, see Checking Whether the max_wal_senders Value of the Source Database Is Correctly Configured.
		The replica identity attribute of tables that do not have primary keys in the source database must be full .	-

Category	Check Item	Check Item Details	Solution to Failure
		If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full.	For details, see Checking Replication Attribute of Primary Key Columns .
		The wal_level value of the source database must be logical .	For details, see Checking Whether the WAL_LEVEL Value in the Source Database Is Correct .
Incremental synchronization	Table fields	<ul style="list-style-type: none"> If there is a table containing fields of the bytea or text type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, an OOM exception may occur. Supported field types: digit, currency, character, binary, date/time, Boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. 	-
	Unlogged tables	During database-level synchronization, DMLs of unlogged tables cannot be synchronized.	-
	Parameters	To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica . After the synchronization is complete, change the value to origin .	For details, see Checking Whether session_replication_role of the Destination Database Is Correctly Set .
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space .
	Status	The destination DB instance is running properly.	-

Category	Check Item	Check Item Details	Solution to Failure
	Character type and sorting rule	The destination database does not support the value of <code>lc_ctype</code> or <code>lc_collate</code> in the database to be synchronized.	For details, see Character Type and Sorting Rule Check in the Destination Database .
	Currency format	The currency format of the source and destination databases must be the same.	For details, see Checking Whether the <code>lc_monetary</code> Values of the Source and Destination Databases Are the Same .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> Database- and table-level synchronization is supported. Instance-level synchronization is not supported. Objects that are not supported: system schemas (schemas starting with <code>pg_</code>, <code>information_schema</code>, <code>sys</code>, <code>utl_raw</code>, <code>dbms_lob</code>, <code>dbms_output</code>, and <code>dbms_random</code>), system catalogs, system users, tablespaces, foreign-data wrappers, foreign servers, user mappings, publications, and subscriptions 	-
		The selected table cannot contain delay constraints.	For details, see Whether the Selected Table Contains Delay Constraints .
		The objects that have dependencies must be synchronized at the same time.	-
	Naming rules of databases, tables, and columns	<ul style="list-style-type: none"> The database name cannot contain the following special characters: <code>+"%'\<></code> The schema name and table name cannot contain the following special characters: <code>":'<></code> The column name cannot contain double quotation marks (<code>"</code>) or single quotation marks (<code>'</code>). 	-

Category	Check Item	Check Item Details	Solution to Failure
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
		<ul style="list-style-type: none"> For a full+incremental synchronization task, the source database cannot be a standby database. For a full synchronization task, the source database can be a standby database, and hot_standby_feedback must be set to on. 	For details, see Checking Whether the Source Database Is on Standby .
		Before starting a task, ensure that the source database does not contain transactions that have been there for a long time without being submitted.	-
		The SSL status of the source database is normal.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

PostgreSQL->GaussDB(DWS)

Table 15-9 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	Full+incremental synchronization requires the following permissions: The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections	-
	Destination database permissions	<ul style="list-style-type: none"> ● Database-level: The CREATEDB permission is required. ● Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize schemas, the CONNECT and CREATE permissions for the database that contains the schemas are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the objects are required. 	-
Versions	Source database versions	Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .
	Destination database versions	Versions 8.1.3 and 8.2.0 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
Parameters	Performance parameters	The max_replication_slots value of the source database must be greater than the number of used replication slots.	For details, see Checking Whether the MAX_REPLICATION_SLOTS Value in the Source Database Is Correct .
		The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value.	For details, see Checking Whether the max_wal_senders Value of the Source Database Is Correctly Configured .
		The replica identity attribute of tables that do not have primary keys in the source database must be full .	-
		If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full.	For details, see Checking Replication Attribute of Primary Key Columns .
		The wal_level value of the source database must be logical .	For details, see Checking Whether the WAL_LEVEL Value in the Source Database Is Correct .
Incremental synchronization	Table fields	<ul style="list-style-type: none"> • GaussDB(DWS)-compatible PostgreSQL data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz and interval, are supported. • XML, line, and domain data types cannot be synchronized. 	-
	Unlogged tables	During database-level synchronization, DMLs of unlogged tables cannot be synchronized.	-

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Schemas, tables, indexes, constraints, sequences, and user-defined types can be synchronized. • During table-level synchronization, only tables and sequences can be synchronized. • Instance-level synchronization is not supported. • System schemas (any schema starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random) and system catalogs cannot be synchronized. System schemas (any schema starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random) and system catalogs cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. 	-
		The selected table cannot contain delay constraints.	For details, see Whether the Selected Table Contains Delay Constraints.
		The objects that have dependencies must be synchronized at the same time. Common dependencies are as follows: tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.	-
	Naming rules of databases, tables, and columns	<ul style="list-style-type: none"> • The database name cannot contain the following special characters: +'%'\<> • The schema name and table name cannot contain the following special characters: ".'\<> • The column name cannot contain double quotation marks (") or single quotation marks ('). 	-

Category	Check Item	Check Item Details	Solution to Failure
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
		<ul style="list-style-type: none"> For an incremental synchronization task, the source database cannot be a standby database. For a full synchronization task, the source database can be a standby database, and hot_standby_feedback must be set to on. 	For details, see Checking Whether the Source Database Is on Standby .
		Before starting a task, ensure that the source database does not contain transactions that have been there for a long time without being submitted.	-
		The SSL status of the source database is normal.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

PostgreSQL -> GaussDB Primary/Standby

Table 15-10 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following permissions: The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences Full+incremental synchronization requires the following permissions: The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections 	-
	Destination database permissions	The user has the sysadmin role or the following minimum permissions: The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables.	-
Versions	Source database versions	Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Parameters	Performance parameters	The max_replication_slots value of the source database must be greater than the number of used replication slots.	For details, see Checking Whether the MAX_REPLICATION_SLOTS Value in the Source Database Is Correct .

Category	Check Item	Check Item Details	Solution to Failure
		The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value.	For details, see Checking Whether the max_wal_senders Value of the Source Database Is Correctly Configured.
		The replica identity attribute of tables that do not have primary keys in the source database must be full .	-
		If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full.	For details, see Checking Replication Attribute of Primary Key Columns.
		The wal_level value of the source database must be logical .	For details, see Checking Whether the WAL_LEVEL Value in the Source Database Is Correct.
Incremental synchronization	Table fields	<ul style="list-style-type: none"> • GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz and interval, are supported. • XML, line, domain, and self-built data type synchronization is not supported. 	-
	Unlogged tables	During database-level synchronization, DMLs of unlogged tables cannot be synchronized.	-
Consistency	Schemas and tables	The schemas and tables of the destination database must be the same as those of the source database.	-

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Schemas, tables, primary key and unique constraints, table data, and sequences can be synchronized. • Instance-level synchronization is not supported. • System schemas and system catalogs cannot be synchronized. (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. 	-
		The selected table cannot contain delay constraints.	For details, see Whether the Selected Table Contains Delay Constraints .
		The objects that have dependencies must be synchronized at the same time. Common dependencies are as follows: tables referenced by views, views referenced by views, tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.	-

Category	Check Item	Check Item Details	Solution to Failure
	Naming rules of databases, tables, and columns	<ul style="list-style-type: none"> The database name cannot contain the following special characters: +"%'\<> The schema name and table name cannot contain the following special characters: ".'\<> The column name cannot contain double quotation marks (") or single quotation marks ('). 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
		<ul style="list-style-type: none"> For an incremental synchronization task, the source database cannot be a standby database. For a full synchronization task, the source database can be a standby database, and hot_standby_feedback must be set to on. 	For details, see Checking Whether the Source Database Is on Standby .
		Before starting a task, ensure that the source database does not contain transactions that have been there for a long time without being submitted.	-
		The SSL status of the source database is normal.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

PostgreSQL -> GaussDB Distributed

Table 15-11 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following permissions: The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, and the SELECT permission for sequences Full+incremental synchronization requires the following permissions: The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections 	-
	Destination database permissions	The user has the sysadmin role or the following minimum permissions: The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables.	-
Versions	Source database versions	Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Parameters	Performance parameters	The max_replication_slots value of the source database must be greater than the number of used replication slots.	For details, see Checking Whether the MAX_REPLICATION_SLOTS Value in the Source Database Is Correct .

Category	Check Item	Check Item Details	Solution to Failure
		The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value.	For details, see Checking Whether the max_wal_senders Value of the Source Database Is Correctly Configured.
		The replica identity attribute of tables that do not have primary keys in the source database must be full .	-
		If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full.	For details, see Checking Replication Attribute of Primary Key Columns.
		The wal_level value of the source database must be logical .	For details, see Checking Whether the WAL_LEVEL Value in the Source Database Is Correct.
Incremental synchronization	Table fields	<ul style="list-style-type: none"> GaussDB (PostgreSQL-compatible) data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz and interval, are supported. XML, line, domain, and self-built data type synchronization is not supported. 	-
	Unlogged tables	During database-level synchronization, DMLs of unlogged tables cannot be synchronized.	-
Consistency	Schemas and tables	The schemas and tables of the destination database must be the same as those of the source database.	-

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Schemas, tables, primary key and unique constraints, table data, and sequences can be synchronized. • Instance-level synchronization is not supported. • System schemas and system catalogs cannot be synchronized. (The following schemas are system schemas in the destination database and cannot be synchronized: "pg_toast", "cstore", "snapshot", "sys", "dbms_job", "dbms_perf", "pg_catalog", "information_schema", "utl_file", "dbms_output", "dbms_random", "utl_raw", "dbms_sql", "dbms_lob", "dbe_perf", "pkg_service", "pkg_util", "dbe_file", "dbe_random", "dbe_output", "dbe_raw", "dbe_sql", "dbe_lob", "dbe_task", "blockchain", "db4ai", "dbe_pldebugger", "sqladvisor", "dbe_application_info", "dbe_match", "dbe_pldeveloper", "dbe_scheduler", "dbe_session", "dbe_utility", "dbe_sql_util"). • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. 	-
		The selected table cannot contain delay constraints.	For details, see Whether the Selected Table Contains Delay Constraints .
		The objects that have dependencies must be synchronized at the same time. Common dependencies are as follows: tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.	-

Category	Check Item	Check Item Details	Solution to Failure
	Naming rules of databases, tables, and columns	<ul style="list-style-type: none"> The database name cannot contain the following special characters: + "%' \ < > The schema name and table name cannot contain the following special characters: ". ' < > The column name cannot contain double quotation marks (") or single quotation marks ('). 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
		<ul style="list-style-type: none"> For an incremental synchronization task, the source database cannot be a standby database. For a full synchronization task, the source database can be a standby database, and hot_standby_feedback must be set to on. 	For details, see Checking Whether the Source Database Is on Standby .
		Before starting a task, ensure that the source database does not contain transactions that have been there for a long time without being submitted.	-
		The SSL status of the source database is normal.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

Oracle->MySQL

Table 15-12 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization:</p> <p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+incremental synchronization and incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. 	
	Destination database permissions	The user must have the SELECT, INSERT, CREATE, DROP, UPDATE, ALTER, DELETE and INDEX permissions.	-
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
Incremental synchronization	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	Compatibility with heterogeneous databases	<ul style="list-style-type: none"> The primary key or unique key column cannot contain data of the character string type when you map the MySQL data types to the character string data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur. Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle. 	For details, see Checking Whether the Values in the Source Oracle Database Are Out of the Ranges of the MySQL Database.
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns.
	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible.
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Status	The destination DB instance is running properly.	-
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> Table-level synchronization or object file import is supported. Databases, table structures, primary keys, unique keys, normal indexes, and table data can be synchronized. Other database objects, such as stored procedures, triggers, functions, sequences, packages, synonyms, and users, cannot be synchronized. 	-
	Table structures	The table structure of the source database must be the same as that of the destination database.	-
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Collation check	The primary key or unique key in the destination database contains a collation ending with _ci. The collation ending with _ci is case insensitive, so an error indicating duplicate keys may be reported during synchronization.	For details, see Checking Whether the Collation of the Destination Database Is Correct .

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints.
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported.
	Duplicate names	Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.
	Database table naming rules	<ul style="list-style-type: none"> Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\ ?!" The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.

Category	Check Item	Check Item Details	Solution to Failure
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

Oracle->GaussDB(for MySQL)

Table 15-13 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+Incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs. 	
	Destination database permissions	The user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD, CREATE VIEW, CREATE ROUTINE and TRIGGER.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	Version 8.0 is supported.	For details, see Supported Databases .
Incremental synchronization	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	Compatibility with heterogeneous databases	<ul style="list-style-type: none"> The primary key or unique key column cannot contain data of the character string type when you map the MySQL data types to the character string data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur. Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle. 	For details, see Checking Whether the Values in the Source Oracle Database Are Out of the Ranges of the MySQL Database.
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns.
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-
Destination database	Status	The destination DB instance is running properly.	-

Category	Check Item	Check Item Details	Solution to Failure
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database .
Synchronization objects	Selected objects	Databases, table structures, primary keys, unique keys, normal indexes, and table data can be synchronized. Other database objects, such as stored procedures, triggers, functions, sequences, packages, synonyms, and users, cannot be synchronized.	-
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Collation check	The primary key or unique key in the destination database contains a collation ending with _ci. The collation ending with _ci is case insensitive, so an error indicating duplicate keys may be reported during synchronization.	For details, see Checking Whether the Collation of the Destination Database Is Correct .
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints .

Category	Check Item	Check Item Details	Solution to Failure
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported.
	Duplicate names	Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.
	Database table naming rules	<ul style="list-style-type: none"> The names of databases and tables cannot contain non-ASCII characters or the following special characters: .>`<'\" ?!" The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Oracle -> GaussDB Primary/Standby

Table 15-14 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization:</p> <p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+incremental synchronization and incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none">• To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.	

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If <code>gs_loader</code> is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-

Category	Check Item	Check Item Details	Solution to Failure
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Incremental synchronization	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible .
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns .
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-

Category	Check Item	Check Item Details	Solution to Failure
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. • Tables can be synchronized in real time during incremental synchronization. 	-
	Whether the selected objects exist in the destination database	Check whether the destination database objects meet the synchronization requirements.	For details, see Whether the Selected Objects Exist in the Destination Database .
	Foreign keys	Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.	For details, see Whether Foreign Keys Are Disabled or Tables to Be Synchronized Have Foreign Keys in the Destination Database .
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints.
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported.
	Database table naming rules	Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Oracle -> GaussDB Distributed

Table 15-15 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization:</p> <p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+incremental synchronization and incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> • To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. 	

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> If <code>gs_loader</code> is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-

Category	Check Item	Check Item Details	Solution to Failure
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Incremental synchronization	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible .
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns .
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-

Category	Check Item	Check Item Details	Solution to Failure
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. • Tables can be synchronized in real time during incremental synchronization. 	-
	Whether the selected objects exist in the destination database	Check whether the destination database objects meet the synchronization requirements.	For details, see Whether the Selected Objects Exist in the Destination Database .
	Foreign keys	Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.	For details, see Whether Foreign Keys Are Disabled or Tables to Be Synchronized Have Foreign Keys in the Destination Database .
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints.
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported.
	Database table naming rules	Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Oracle->GaussDB(DWS)

Table 15-16 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization:</p> <p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+incremental synchronization and incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. 	
	Destination database permissions	<ul style="list-style-type: none"> If the destination side does not contain databases, schemas, or tables, the destination database user must have the permission to create databases, the permission to create schemas in a database, or the permission to create tables in a schema. The INSERT, SELECT, UPDATE, and DELETE permissions are required for each table. 	-
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	Versions 8.1.3 and 8.2.0 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
Incremental synchronization	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible.
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns.
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-
Destination database	Status	The destination DB instance is running properly.	-
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, common indexes (B-Tree indexes), and constraints (primary key, empty, and not-null) can be synchronized. • Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, unique indexes, foreign key indexes, and check constraints cannot be synchronized. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Duplicate names	During table structure synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same .
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints .
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported .
	Database table naming rules	<ul style="list-style-type: none"> The names of databases and tables cannot contain non-ASCII characters or the following special characters: .><\,?!" An empty source database cannot be synchronized. 	-

Category	Check Item	Check Item Details	Solution to Failure
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Oracle->DDM

Table 15-17 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization: The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+Incremental synchronization:</p> <ul style="list-style-type: none"> ● Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. ● Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. 	
	Destination database permissions	The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD and CREATE VIEW	-
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
Incremental synchronization	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible .
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns .

Category	Check Item	Check Item Details	Solution to Failure
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database .
Synchronization objects	Selected objects	The source database data can be synchronized, but the source table structure and other objects cannot be synchronized.	-
	Table structures	The redundant columns (columns that do not exist in the source database) of the destination database cannot contain not-null constraints. The not-null constraints will cause the migration to fail.	For details, see Whether the Table Structures (Including Primary Key Indexes and the Number of Columns) of the Source Oracle Database and Destination Database Middleware Are Aligned .

Category	Check Item	Check Item Details	Solution to Failure
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints .
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • Tables whose names contain special characters except underscores (_) cannot be synchronized. • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • An empty source database cannot be synchronized. • Before synchronization, you must create databases, tables, columns, indexes, and constraints with the same names (in lowercase) as the corresponding source objects to be synchronized in the destination database. 	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

Oracle->PostgreSQL

Table 15-18 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<p>Full synchronization:</p> <p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser)</p> <p>Full+Incremental synchronization:</p> <ul style="list-style-type: none"> • Oracle 12c or later in tenant mode: <ul style="list-style-type: none"> - To synchronize a container database (CDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. • Oracle 12c or later in non-tenant mode: 	-

Category	Check Item	Check Item Details	Solution to Failure
		<p>The user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <ul style="list-style-type: none"> To synchronize a database of Oracle 11g or earlier, the user must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. 	
	Destination database permissions	The user must have the following permissions for a table: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.	-
	MAINTAINED check of the source Oracle database user	Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permissions to parse logs.	-
Versions	Source database versions	Versions 10g, 11g, 12c, 18c, 19c, and 21c are supported.	For details, see Supported Databases .
	Destination database versions	Versions 9.5, 9.6, 10, 11, 12, 13, and 14 are supported.	For details, see Supported Databases .
Incremental synchronization	Compatibility with heterogeneous databases	Ensure that the columns that use character strings as primary keys or unique keys in the table to be migrated do not contain spaces. Otherwise, data may be inconsistent or the migration may fail.	For details, see Checking Whether the Values in the Source Oracle Database Are Out of the Ranges of the MySQL Database .

Category	Check Item	Check Item Details	Solution to Failure
	Virtual columns	The source database cannot contain virtual columns.	For details, see Checking Whether the Source Table Structure Contains Virtual Columns.
	Character set compatibility check	The character set of the destination database must be fully compatible with that of the source database.	For details, see Checking Whether the Character Sets Are Compatible.
	Archive logs and supplementary logs	During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.	-
Destination database	Status	The destination DB instance is running properly.	-
	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
Source database	Physical standby database check	During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.	For details, see Checking the Physical Standby Database.

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, indexes, and constraints (primary key, null, and non-null) can be synchronized. • Views, foreign keys, stored procedures, triggers, functions, events, and virtual columns cannot be synchronized. 	-
	Table field types	Check whether the source database contains unsupported table field types. The following table field types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, RAW, LONG RAW, CHAR, NCHAR, CLOB, NCLOB, BLOB, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE.	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Data constraints	During data transmission, data written to the destination database may fail or be lost because the data fails the constraint check.	For details, see Checking Whether Existing Data Meets the Constraints .
	Character sets	Only the following character sets are supported for the source database: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.	For details, see Checking Whether the Source Database Character Set Is Supported .
	Log parsing permissions	If the source database user does not have the log parsing permission, the incremental migration will fail.	For details, see Checking Whether the Source Database User Has the Permission to Parse Logs .

Category	Check Item	Check Item Details	Solution to Failure
	Duplicate names	Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same.
	Database table naming rules	The names of databases and tables cannot contain non-ASCII characters or the following special characters: .>< \,;?!"	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

DDM->MySQL

Table 15-19 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The user of the source DDM database must have at least one permission, for example, SELECT. The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	The user must have the SELECT, CREATE, ALTER, DROP, DELETE, INSERT and UPDATE permissions. The root account of the RDS for MySQL DB instance has the preceding permissions by default.	For details, see Which MySQL Permissions Are Required for DRS?
Parameters	GTID status	Enable GTID for the source sharded database. If GTID is not enabled for the source sharded database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	-
	Performance parameters	The binlog_row_image parameter of the source sharded database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL.
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source sharded database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled.
	Binlog format	The source sharded database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based.
	server_id value	During an incremental synchronization, the server_id value of the source sharded database must be specified. <ul style="list-style-type: none"> If the source sharded database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. If the source sharded database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
Synchronization objects	Selected objects	<ul style="list-style-type: none"> Full synchronization supports the synchronization of data, table structures, and indexes. The source database cannot contain tables whose sharding keys are of the timestamp type. 	-

Category	Check Item	Check Item Details	Solution to Failure
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Database table naming rules	The database names and table names of the source sharding middleware cannot contain non-ASCII characters and the following characters: '<>\'	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured .
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

DDM->GaussDB(DWS)

Table 15-20 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The user of the source DDM database must have at least one permission, for example, SELECT. The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions. 	-
	Destination database permissions	The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE.	-
Parameters	GTID status	Enable GTID for the source sharded database. If GTID is not enabled for the source sharded database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	-
	Performance parameters	The binlog_row_image parameter of the source sharded database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL.
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source sharded database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled.
	Binlog format	The source sharded database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based.

Category	Check Item	Check Item Details	Solution to Failure
	server_id value	<p>During an incremental synchronization, the server_id value of the source sharded database must be specified.</p> <ul style="list-style-type: none"> • If the source sharded database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source sharded database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	<p>For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.</p>
Destination database	Storage space	<p>The destination DB instance must have sufficient storage space.</p>	<p>For details, see Checking Whether the Destination Database Has Sufficient Storage Space.</p>
Consistency	Character set	<p>The character set of the destination database must be the same as that of the source database.</p>	<p>For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.</p>
	Case sensitive	<p>The lower_case_table_names value of the destination database must be the same as that of the source database.</p>	<p>For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.</p>
	Maximum length of the calculation result of the group_concat function	<p>The group_concat_max_len value of the destination database must be the same as that of the source database.</p>	<p>For details, see Checking Whether the Values of group_concat_max_len Are Consistent.</p>

Category	Check Item	Check Item Details	Solution to Failure
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Full synchronization supports the synchronization of data, table structures, and common indexes (B-Tree indexes). • The destination database does not support tables with unique keys. During synchronization, the tables with unique keys in the source database are ignored. • A table without a primary key cannot be synchronized. If the selected table does not have a primary key, the synchronization fails. • Unique constraints and indexes cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	-
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	The database names and table names of the source sharding middleware cannot contain non-ASCII characters and the following characters: '<>\'	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

DDM->DDM

Table 15-21 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The user of the source DDM database must have at least one permission, for example, SELECT. The DDM sharded database user must have the SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT permissions. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	The user must have the following permissions: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT. In addition, grant the SELECT permission on all tables. The DDM destination database user must have permissions on the database to be synchronized.	-
Parameters	GTID status	Enable GTID for the source sharded database. If GTID is not enabled for the source sharded database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	-
	Performance parameters	The binlog_row_image parameter of the source sharded database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source sharded database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source sharded database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .
	server_id value	During an incremental synchronization, the server_id value of the source sharded database must be specified. <ul style="list-style-type: none"> • If the source sharded database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source sharded database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Full synchronization supports the synchronization of data, table structures, and indexes. • Only MyISAM and InnoDB tables can be synchronized. • The source database cannot contain tables whose sharding keys are of the timestamp type. 	-
		During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects.	For details, see Checking Database Mapping Objects .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Database table naming rules	The database names and table names of the source sharding middleware cannot contain non-ASCII characters and the following characters: '<>\'	-
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured .
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

DB2 for LUW -> GaussDB Primary/Standby

Table 15-22 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The CONNECT and DATAACCESS permissions are mandatory for full synchronization. The DBADM permission is mandatory for full+incremental synchronization. If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table. 	-
	Destination database permissions	<ul style="list-style-type: none"> Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. Schema-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the CREATE and USAGE permissions on schemas to the user. Table-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the SELECT, UPDATE, INSERT, and DELETE permissions on all tables in the schema to the user. If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-
Versions	Source database versions	Versions 9.7, 10.1, 10.5, 11.1, and 11.5 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Parameters	DataCapture	If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the DataCapture attribute of the table to be synchronized is Y.	-
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. • Tables can be synchronized in real time during incremental synchronization. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. 	-
	Whether the selected objects exist in the destination database	Check whether the destination database objects meet the synchronization requirements.	For details, see Whether the Selected Objects Exist in the Destination Database .

Category	Check Item	Check Item Details	Solution to Failure
	Table field types	<ul style="list-style-type: none"> • The maximum precision supported by the TIMESTAMP type is 6. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • User-defined types are not supported. • VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. 	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Whether the Source Database Tables Contain Primary Keys .
	Character sets	The source database supports only the GBK and UTF8 character sets.	For details, see Checking Whether the Source Database Character Set Is Supported .
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. • Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. 	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

DB2 for LUW -> GaussDB Distributed

Table 15-23 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The CONNECT and DATAACCESS permissions are mandatory for full synchronization. The DBADM permission is mandatory for full+incremental synchronization. If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table. 	-
	Destination database permissions	<ul style="list-style-type: none"> Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. Schema-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the CREATE and USAGE permissions on schemas to the user. Table-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the SELECT, UPDATE, INSERT, and DELETE permissions on all tables in the schema to the user. If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-
Versions	Source database versions	Versions 9.7, 10.1, 10.5, 11.1, and 11.5 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	1.0.0 and later versions are supported.	For details, see Supported Databases .
Parameters	DataCapture	If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the DataCapture attribute of the table to be synchronized is Y.	-
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. • Tables can be synchronized in real time during incremental synchronization. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. 	-
	Whether the selected objects exist in the destination database	Check whether the destination database objects meet the synchronization requirements.	For details, see Whether the Selected Objects Exist in the Destination Database .

Category	Check Item	Check Item Details	Solution to Failure
	Table field types	<ul style="list-style-type: none"> • The maximum precision supported by the TIMESTAMP type is 6. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • User-defined types are not supported. • VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. 	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Whether the Source Database Tables Contain Primary Keys .
	Character sets	The source database supports only the GBK and UTF8 character sets.	For details, see Checking Whether the Source Database Character Set Is Supported .
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • During table structure synchronization in the full phase, if there is a schema with the same name as a user in the destination database and another user is used to synchronize the table structure to the schema, run the grant [role] to [role] command to grant permissions to the user. Otherwise, the synchronization may fail because the table structure cannot be created. • Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. 	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

DB2 for LUW->GaussDB(DWS)

Table 15-24 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> The CONNECT and DATAACCESS permissions are mandatory for full synchronization. The DBADM permission is mandatory for full+incremental synchronization. If there is the DB2SECURITYLABEL data in the table structure of the source database, ensure that the user has the read permission on all data in the table. 	-
	Destination database permissions	<ul style="list-style-type: none"> If the destination side does not contain databases, schemas, or tables, the destination database user must have the permission to create databases, the permission to create schemas in a database, or the permission to create tables in a schema. The INSERT, SELECT, UPDATE, and DELETE permissions are required for each table. 	-
Versions	Source database versions	Versions 9.7, 10.1, 10.5, 11.1, and 11.5 are supported.	For details, see Supported Databases .
	Destination database versions	Versions 8.1.3 and 8.2.0 are supported.	For details, see Supported Databases .
Parameters	DataCapture	If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the DataCapture attribute of the table to be synchronized is Y.	-

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary key and unique constraints, and data can be synchronized. • Tables can be synchronized in real time during incremental synchronization. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. • The number of synchronization objects of a single schema in the source database cannot exceed 32,766. 	-
	Whether the selected objects exist in the destination database	Check whether the destination database objects meet the synchronization requirements.	For details, see Whether the Selected Objects Exist in the Destination Database .

Category	Check Item	Check Item Details	Solution to Failure
	Table field types	<ul style="list-style-type: none"> • The maximum precision supported by the TIMESTAMP type is 6. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • User-defined types are not supported. • VARCHAR or VARGRAPHIC data whose length exceeds 3998 characters cannot be synchronized during the incremental synchronization. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	For details, see Checking Whether the Source Database Contains Unsupported Table Field Types .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Whether the Source Database Tables Contain Primary Keys .
	Character sets	The source database supports only the GBK and UTF8 character sets.	For details, see Checking Whether the Source Database Character Set Is Supported .
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or the following special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.	-
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

TiDB->GaussDB(for MySQL)

Table 15-25 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	The user must have the SELECT and CONFIG permissions.	-
	Destination database permissions	The user must have the SELECT, CREATE, DROP, INSERT, DELETE, UPDATE, ALTER, REFERENCES and INDEX permissions.	-
Versions	Source database versions	Version 4.0.0 and later (excluding the development version)	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	Version 8.0	For details, see Supported Databases .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> Table structures, data, and indexes and constraints of selected tables can be synchronized. Supported field types BIGINT, BINARY, BIT, BLOB, BOOLEAN, CHAR, DATE, DATETIME, DECIMAL, DOUBLE, ENUM, FLOAT, INT, JSON, LONGBLOB, LONGTEXT, MEDIUMBLOB, MEDIUMINT, SET, SMALLINT, TEXT, TIME, TIMESTAMP, TINYBLOB, TINYINT, TINYTEXT, VARBINARY, VARCHAR, YEAR 	-
	Database table naming rules	<ul style="list-style-type: none"> The database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. The names of the databases and tables to be synchronized cannot contain non-ASCII characters or the following special characters: '<'>\/. 	-
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Whether the Source Database Tables Contain Primary Keys .

Category	Check Item	Check Item Details	Solution to Failure
	Duplicate names	The destination database cannot contain tables with the same names as those in the source database.	For details, see Checking Whether Destination Database Contains the Same Table Names As the Synchronization Objects.
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Microsoft SQL Server->GaussDB(DWS)

Table 15-26 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> • If the destination instance does not contain the database to be synchronized, the CREATEDB permission is required. • If the destination instance contains databases but does not have schemas, the CONNECT and CREATE permissions for the databases are required. • If the destination instance contains databases and schemas but does not contain tables, the CONNECT permission for databases and the USAGE and CREATE permissions for schemas are required. • If the destination instance contains databases, schemas, and tables, the CONNECT permission for databases, the USAGE permission for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, REFERENCES, and INDEX permissions for tables are required. 	-
Versions	Source database versions	<ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) 	For details, see Supported Databases .
	Destination database versions	Versions 8.1.3 and 8.2.0 are supported.	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
Incremental synchronization	CDC enabling	Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.	For details, see Checking Whether CDC is Enabled for Tables in the Source Database .
	CDC data retention	If the CDC data is retained less than one day in the source database, incremental synchronization will be abnormal. Change the retention period to 1440 minutes (one day) or longer. The recommended value is 4320 minutes (three days).	For details, see Checking Whether the CDC Retention Period in the Source Database Is Long Enough .
Consistency	Computer names	The computer names of the source and destination databases must be different.	For details, see Checking Whether the Source and Destination Databases Have Different Computer Names .

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Table structures, data, and indexes of selected tables can be synchronized. ● The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) ● Column encryption is not supported. ● Auto-increment columns cannot be synchronized. ● Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY 	-
	Database table naming rules	<ul style="list-style-type: none"> ● Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). ● Names of the columns in the source table cannot contain the following special characters: []? ● The source database cannot contain the username cdc or schema. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Duplicate names	The destination database cannot contain objects with the same name as those in the source database (except SQL Server system databases).	For details, see Checking Whether the Destination Database Contains Objects with the Same Name As Those in the Source Database .
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects .
		The source database cannot be empty.	For details, see Checking Whether the Source Database Is Empty .
		The recovery model of the source database must be set to FULL.	For details, see Checking Whether the Source Database Uses the Full Recovery Model .
		SQL Server Agent must be enabled for the source database.	For details, see Checking Whether the SQL Server Agent of the Source Database Is Enabled .

Category	Check Item	Check Item Details	Solution to Failure
		If the source database contains disabled clustered indexes of tables, the synchronization fails.	For details, see Checking Whether the Source Database Contains Disabled Clustered Indexes.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Microsoft SQL Server -> GaussDB Primary/Standby

Table 15-27 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following minimum permissions: Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized Full+incremental synchronization and incremental synchronization require the following minimum permissions: Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized 	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> ● Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. ● Schema-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the CREATE and USAGE permissions on schemas to the user. ● Table-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the SELECT, UPDATE, INSERT, and DELETE permissions on all tables in the schema to the user. ● If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-
Versions	Source database versions	<ul style="list-style-type: none"> ● On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) 	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	Version 1.1.0 or later	For details, see Supported Databases .
Incremental synchronization	CDC enabling	Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.	For details, see Checking Whether CDC is Enabled for Tables in the Source Database .
	CDC data retention	If the CDC data is retained less than one day in the source database, incremental synchronization will be abnormal. Change the retention period to 1440 minutes (one day) or longer. The recommended value is 4320 minutes (three days).	For details, see Checking Whether the CDC Retention Period in the Source Database Is Long Enough .
Consistency	Computer names	The computer names of the source and destination databases must be different.	For details, see Checking Whether the Source and Destination Databases Have Different Computer Names .

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Table structures, data, and indexes of selected tables can be synchronized. ● The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) ● Column encryption is not supported. ● Auto-increment columns cannot be synchronized. ● Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY 	-
	Database table naming rules	<ul style="list-style-type: none"> ● Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). ● Names of the columns in the source table cannot contain the following special characters: []? ● The source database cannot contain the username cdc or schema. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Duplicate names	The destination database cannot contain objects with the same name as those in the source database (except SQL Server system databases).	For details, see Checking Whether the Destination Database Contains Objects with the Same Name As Those in the Source Database .
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects .
		The source database cannot be empty.	For details, see Checking Whether the Source Database Is Empty .
		The recovery model of the source database must be set to FULL.	For details, see Checking Whether the Source Database Uses the Full Recovery Model .
		SQL Server Agent must be enabled for the source database.	For details, see Checking Whether the SQL Server Agent of the Source Database Is Enabled .

Category	Check Item	Check Item Details	Solution to Failure
		If the source database contains disabled clustered indexes of tables, the synchronization fails.	For details, see Checking Whether the Source Database Contains Disabled Clustered Indexes.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Microsoft SQL Server -> GaussDB Distributed

Table 15-28 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full synchronization requires the following minimum permissions: Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized Full+incremental synchronization and incremental synchronization require the following minimum permissions: Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized 	-

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> ● Database-level permissions: Log in to a Postgres database as the root user or the DATABASE user with the sysadmin role. Grant the CREATE and CONNECT permissions to the DATABASE user. ● Schema-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the CREATE and USAGE permissions on schemas to the user. ● Table-level permissions: Log in to the database as the root user, the DATABASE user with the sysadmin role, or the owner of the database. Grant the SELECT, UPDATE, INSERT, and DELETE permissions on all tables in the schema to the user. ● If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. 	-
Versions	Source database versions	<ul style="list-style-type: none"> ● On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ● RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) 	For details, see Supported Databases .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	Version 1.1.0 or later	For details, see Supported Databases .
Incremental synchronization	CDC enabling	Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.	For details, see Checking Whether CDC is Enabled for Tables in the Source Database .
	CDC data retention	If the CDC data is retained less than one day in the source database, incremental synchronization will be abnormal. Change the retention period to 1440 minutes (one day) or longer. The recommended value is 4320 minutes (three days).	For details, see Checking Whether the CDC Retention Period in the Source Database Is Long Enough .
Consistency	Computer names	The computer names of the source and destination databases must be different.	For details, see Checking Whether the Source and Destination Databases Have Different Computer Names .

Category	Check Item	Check Item Details	Solution to Failure
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Table structures, data, and indexes of selected tables can be synchronized. ● The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) ● Column encryption is not supported. ● Auto-increment columns cannot be synchronized. ● Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER ● Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY 	-
	Database table naming rules	<ul style="list-style-type: none"> ● Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). ● Names of the columns in the source table cannot contain the following special characters: []? ● The source database cannot contain the username cdc or schema. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Duplicate names	The destination database cannot contain objects with the same name as those in the source database (except SQL Server system databases).	For details, see Checking Whether the Destination Database Contains Objects with the Same Name As Those in the Source Database .
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects .
		The source database cannot be empty.	For details, see Checking Whether the Source Database Is Empty .
		The recovery model of the source database must be set to FULL.	For details, see Checking Whether the Source Database Uses the Full Recovery Model .
		SQL Server Agent must be enabled for the source database.	For details, see Checking Whether the SQL Server Agent of the Source Database Is Enabled .

Category	Check Item	Check Item Details	Solution to Failure
		If the source database contains disabled clustered indexes of tables, the synchronization fails.	For details, see Checking Whether the Source Database Contains Disabled Clustered Indexes.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

Microsoft SQL Server->Microsoft SQL Server

Table 15-29 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	At least the sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized	-
	Destination database permissions	If the destination end does not contain databases, the destination database user must have the create any database permission. If the destination end contains databases, the destination database user must have the connect, create table, alter any schema, and select permissions for the databases.	-

Category	Check Item	Check Item Details	Solution to Failure
Versions	Source database versions	<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) 	For details, see Supported Databases .
	Destination database versions	RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019)	For details, see Supported Databases .
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Destination database	Status	The destination DB instance is running properly.	-
Incremental synchronization	CDC enabling	Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected.	For details, see Checking Whether CDC is Enabled for Tables in the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	CDC data retention	If the CDC data is retained less than one day in the source database, incremental synchronization will be abnormal. Change the retention period to 1440 minutes (one day) or longer. The recommended value is 4320 minutes (three days).	For details, see Checking Whether the CDC Retention Period in the Source Database Is Long Enough .
Consistency	Computer names	The computer names of the source and destination databases must be different.	For details, see Checking Whether the Source and Destination Databases Have Different Computer Names .
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Table structures, data, and indexes of selected tables can be synchronized. • The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) • Column encryption is not supported. • Auto-increment columns cannot be synchronized. • Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER • Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY 	-

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). • Names of the columns in the source table cannot contain the following special characters: []? • The source database cannot contain the username cdc or schema. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. 	-
	Duplicate names	The destination database cannot contain objects with the same name as those in the source database (except SQL Server system databases).	For details, see Checking Whether the Destination Database Contains Objects with the Same Name As Those in the Source Database .
	Source database	The selected synchronization object must exist in the source database.	For details, see Checking the Synchronization Objects .
		The source database cannot be empty.	For details, see Checking Whether the Source Database Is Empty .
		The recovery model of the source database must be set to FULL.	For details, see Checking Whether the Source Database Uses the Full Recovery Model .

Category	Check Item	Check Item Details	Solution to Failure
		SQL Server Agent must be enabled for the source database.	For details, see Checking Whether the SQL Server Agent of the Source Database Is Enabled.
		Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them from the synchronization objects. If you need to synchronize TDE encrypted databases, disable TDE first.	For details, see Whether the Source Database Contains Encrypted Objects.
		If the source database contains disabled clustered indexes of tables, the synchronization fails.	For details, see Checking Whether the Source Database Contains Disabled Clustered Indexes.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MongoDB->DDS

Table 15-30 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	Replica set: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database.	-
	Destination database permissions	The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database.	-
Versions	Source database versions	<ul style="list-style-type: none"> Self-managed MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) DDS (versions 3.2, 3.4, 4.0, 4.2 and 4.4) <p>NOTE The source database cannot be a GeminiDB Mongo instance.</p>	For details, see Supported Databases .
	Destination database versions	DDS (versions 3.4, 4.0, 4.2 and 4.4)	For details, see Supported Databases .
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space .

Category	Check Item	Check Item Details	Solution to Failure
	Status	The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.	-
	Maximum number of chunks	Check whether the maximum number of chunks in the destination database is sufficient to support sharding and splitting of the source database. If the maximum number of chunks is reached, chunks are not split and the write performance is negatively affected.	For details, see Checking Whether the Maximum Number of Chunks in the Destination Database Is Sufficient.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> ● Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. ● System databases (such as local, admin, and config) and system collections cannot be synchronized. If service data is in a system database, run renameCollection to move the service data to the user database. ● The statement for creating a view cannot contain a regular expression. ● Collections that contain the _id field without indexes are not supported. ● The first parameter of BinData() cannot be 2. ● If the source is a cluster instance of version 4.4, composite hash indexes and composite hash shard keys are not supported. ● If the source is a replica set instance of version 4.4, composite hash indexes are not supported. 	-
	Object dependencies	If the roles referenced by accounts to be migrated are not migrated to the destination database, the migration may fail.	-

Category	Check Item	Check Item Details	Solution to Failure
	Associated objects	Associated objects must be synchronized at the same time to avoid synchronization failure due to missing associated objects. Common dependencies: collections referenced by views, and views referenced by views	For details, see Checking Whether Referenced Tables Are Selected for Migration .
	Database table naming rules	<ul style="list-style-type: none"> The source database name cannot contain \. "\$ or spaces. The collection or view name in the source database cannot start with system. or contain the dollar sign (\$). The mapped synchronization object name can contain 1 to 63 characters. The following characters are not allowed: \. "\$ <>. Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization. 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental synchronization, data consistency cannot be ensured when source and destination databases are in different time zones or clocks.		For details, see Checking Whether the Source Database Collections Contain TTL Indexes .	
The source database cannot be empty.		For details, see Checking Whether the Source Database Is Empty .	
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .

Category	Check Item	Check Item Details	Solution to Failure
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

MariaDB->MariaDB

Table 15-31 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)	-
	Destination database permissions	The root account of RDS for MariaDB has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX	-
Versions	Source database versions	Versions 10.3, 10.4, and 10.5 are supported.	For details, see Supported Databases .
	Destination database versions	Version 10.5 is supported.	For details, see Supported Databases .
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .

Category	Check Item	Check Item Details	Solution to Failure
Parameters	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured.
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL.
	Maximum size of a packet that can be transmitted	During synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small.
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values.
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled.
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based.

Category	Check Item	Check Item Details	Solution to Failure
	Binlog retention period	If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database.
	server_id value	During an incremental synchronization, the server_id parameter of the source MariaDB database must be set to a value ranging from 1 to 4294967296.	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.
	Table fields	If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-
	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database.
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.	-

Category	Check Item	Check Item Details	Solution to Failure
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. 	-
		<ul style="list-style-type: none"> • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects. 	For details, see Checking Database Mapping Objects.
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration.

Category	Check Item	Check Item Details	Solution to Failure
	Encrypted table check	Whether the source database contains encrypted tables.	For details, see Source Encrypted Table Check .
	Duplicate names	If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether Destination Database Contains the Same Table Names As the Synchronization Objects .
	Database table naming rules	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<'>/'\" • The source table and view names cannot contain non-ASCII characters, or the following characters: '<'>/'\" • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

MariaDB->MySQL

Table 15-32 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)	-
	Destination database permissions	The root account of RDS for MySQL has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.	-
Versions	Source database versions	Versions 10.0, 10.1, 10.2, 10.3, 10.4, and 10.5 are supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
Parameters	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	During synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small .
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .

Category	Check Item	Check Item Details	Solution to Failure
	Binlog retention period	If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database.
	server_id value	During an incremental synchronization, the server_id parameter of the source MariaDB database must be set to a value ranging from 1 to 4294967296.	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.
	Table fields	If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-
	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database.
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.	-

Category	Check Item	Check Item Details	Solution to Failure
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.

Category	Check Item	Check Item Details	Solution to Failure
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. 	-
		<ul style="list-style-type: none"> • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects. 	For details, see Checking Database Mapping Objects.
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.

Category	Check Item	Check Item Details	Solution to Failure
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration .
	Encrypted table check	Whether the source database contains encrypted tables.	For details, see Source Encrypted Table Check .
	Duplicate names	If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether Destination Database Contains the Same Table Names As the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
	Database table naming rules	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<>/' • The source table and view names cannot contain non-ASCII characters, or the following characters: '<>/' • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured .
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected .

Category	Check Item	Check Item Details	Solution to Failure
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected .

MariaDB->GaussDB(for MySQL)

Table 15-33 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions)	-
	Destination database permissions	The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.	-
Versions	Source database versions	Versions 10.0, 10.1, 10.2, 10.3, 10.4, and 10.5 are supported.	For details, see Supported Databases .
	Destination database versions	Version 8.0 is supported.	For details, see Supported Databases .
Parameters	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .

Category	Check Item	Check Item Details	Solution to Failure
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
	Maximum size of a packet that can be transmitted	During synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small .
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .
	Binlog retention period	If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .

Category	Check Item	Check Item Details	Solution to Failure
	server_id value	During an incremental synchronization, the server_id parameter of the source MariaDB database must be set to a value ranging from 1 to 4294967296.	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.
	Table fields	If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-
	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database.
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.

Category	Check Item	Check Item Details	Solution to Failure
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. 	-
		<ul style="list-style-type: none"> • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects. 	For details, see Checking Database Mapping Objects.
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration.

Category	Check Item	Check Item Details	Solution to Failure
	Encrypted table check	Whether the source database contains encrypted tables.	For details, see Source Encrypted Table Check .
	Duplicate names	If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether Destination Database Contains the Same Table Names As the Synchronization Objects .
	Database table naming rules	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: '<'>/'\" • The source table and view names cannot contain non-ASCII characters, or the following characters: '<'>/'\" • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.

GaussDB(for MySQL)->GaussDB(for MySQL)

Table 15-34 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT	-
	Destination database permissions	The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.	-
Versions	Source database versions	Version 8.0 is supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Destination database versions	Version 8.0 is supported.	For details, see Supported Databases .
	Synchronization version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .
		The binlog_row_image parameter of the source database must be set to FULL . Otherwise, the synchronization will fail.	For details, see Checking Whether the binlog_row_image Value is FULL .
Maximum size of a packet that can be transmitted	During synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small .	

Category	Check Item	Check Item Details	Solution to Failure
		If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the <code>max_allowed_packet</code> value of the source database is too small, the synchronization task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small .
	sql_mode value	If the MyISAM tables are included in the synchronization objects, the <code>sql_mode</code> parameter in the destination database cannot contain the <code>no_engine_substitution</code> parameter. Otherwise, the synchronization fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental synchronization	Binlog status	During the incremental synchronization, the binlog of the source database must be enabled.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	Binlog format	The source database binlog must be row-based.	For details, see Checking Whether the Source Database Binlog Is Row-Based .
	Binlog retention period	If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	During an incremental synchronization, the <code>server_id</code> parameter of the source database must be set to a value ranging from 1 to 4294967296.	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .

Category	Check Item	Check Item Details	Solution to Failure
	Table fields	If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.	-
	Index column length	The index column length of the source database must meet requirements.	For details, see Checking the Length of the Index Column in the Source Database.
Destination database	Status	The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Clock	The clock of the destination database must be the same as that of the source database.	-

Category	Check Item	Check Item Details	Solution to Failure
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
Synchronization objects	Selected objects	<ul style="list-style-type: none"> • Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. • Only MyISAM and InnoDB tables can be synchronized. • Events and triggers cannot be synchronized. • Table names cannot be mapped for tables on which views, stored procedures, and functions depend. 	-
		<ul style="list-style-type: none"> • When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. • During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, or it will result in inconsistent objects. 	For details, see Checking Database Mapping Objects .
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Associated objects	Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys	For details, see Checking Whether Referenced Tables Are Selected for Migration .

Category	Check Item	Check Item Details	Solution to Failure
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database .
	Encrypted table check	Whether the source database contains encrypted tables.	For details, see Source Encrypted Table Check .
	Duplicate names	If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.	For details, see Checking Whether Destination Database Contains the Same Table Names As the Synchronization Objects .
	Database table naming rules	<ul style="list-style-type: none"> • The names of the source databases, tables, and views cannot contain non-ASCII characters, or the following characters: '<>/\'' • The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. • If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. 	-
	Source database	The source database is properly connected during the synchronization object check in the pre-check phase.	For details, see Checking the Synchronization Objects .

Category	Check Item	Check Item Details	Solution to Failure
SSL	SSL configuration	If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.	For details, see Checking Whether the SSL Connection Is Correctly Configured.
Network conditions	Network conditions	The IP address, port, username, and password of the destination database are correctly configured.	For details, see Checking Whether the Destination Database Is Connected.
		The IP address, port, username, and password of the source database are correctly configured.	For details, see Checking Whether the Source Database Is Connected.