

Data Replication Service

Real-Time Synchronization

Issue 45
Date 2023-01-30



Copyright © Huawei Technologies Co., Ltd. 2023. 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://www.huawei.com>

Email: support@huawei.com

Contents

1 Synchronization Overview.....	1
2 Data Synchronization Topologies.....	10
3 To the Cloud.....	13
3.1 From MySQL to MySQL.....	13
3.2 From MySQL to PostgreSQL.....	35
3.3 From MySQL to GaussDB Distributed.....	52
3.4 From MySQL to GaussDB Primary/Standby.....	66
3.5 From MySQL to GaussDB(DWS).....	81
3.6 From MySQL to GaussDB(for MySQL).....	100
3.7 From PostgreSQL to PostgreSQL.....	119
3.8 From PostgreSQL to GaussDB(DWS).....	142
3.9 From PostgreSQL to GaussDB Primary/Standby.....	159
3.10 From PostgreSQL to GaussDB Distributed.....	177
3.11 From Oracle to MySQL.....	195
3.12 From Oracle to GaussDB(for MySQL).....	216
3.13 From Oracle to GaussDB Primary/Standby.....	235
3.14 From Oracle to GaussDB Distributed.....	259
3.15 From Oracle to DDM.....	282
3.16 From Oracle to GaussDB(DWS).....	290
3.17 From Oracle to PostgreSQL.....	308
3.18 From DDM to MySQL.....	326
3.19 From DDM to GaussDB(DWS).....	339
3.20 From DDM to DDM.....	352
3.21 From DB2 for LUW to GaussDB Primary/Standby.....	364
3.22 From DB2 for LUW to GaussDB Distributed.....	378
3.23 From DB2 for LUW to GaussDB(DWS).....	392
3.24 From TiDB to GaussDB(for MySQL).....	406
3.25 From Microsoft SQL Server to GaussDB(DWS).....	422
3.26 From Microsoft SQL Server to GaussDB Primary/Standby.....	438
3.27 From Microsoft SQL Server to GaussDB Distributed.....	453
3.28 From Microsoft SQL Server to Microsoft SQL Server.....	468
3.29 From MongoDB to DDS.....	483

4 Out of the Cloud.....	498
4.1 From MySQL to MySQL.....	498
4.2 From MySQL to Kafka.....	518
4.3 From MySQL to CSS/ES.....	533
4.4 From MySQL to Oracle.....	549
4.5 From DDM to MySQL.....	565
4.6 From DDM to Oracle.....	570
4.7 From DDM to Kafka.....	574
4.8 From DDS to MongoDB.....	585
4.9 From PostgreSQL to Kafka.....	599
4.10 From GaussDB Distributed to MySQL.....	611
4.11 From GaussDB Distributed to Oracle.....	626
4.12 From GaussDB Distributed to GaussDB(DWS).....	646
4.13 From GaussDB Distributed to Kafka.....	665
4.14 From GaussDB Distributed to GaussDB Distributed.....	678
4.15 From GaussDB Distributed to GaussDB Primary/Standby.....	699
4.16 From GaussDB Primary/Standby to MySQL.....	719
4.17 From GaussDB Primary/Standby to Oracle.....	736
4.18 From GaussDB Primary/Standby to Kafka.....	755
4.19 From GaussDB Primary/Standby to GaussDB(DWS).....	769
4.20 From GaussDB Primary/Standby to GaussDB Distributed.....	787
4.21 From GaussDB Primary/Standby to GaussDB Primary/Standby.....	807
4.22 From GaussDB(for MySQL) to MySQL.....	827
4.23 From GaussDB(for MySQL) to GaussDB(DWS).....	842
4.24 From GaussDB(for MySQL) to Kafka.....	858
4.25 From GaussDB(for MySQL) to CSS/ES.....	872
4.26 From GaussDB(for MySQL) to Oracle.....	888
5 Between Self-built Databases.....	903
5.1 From MySQL to Kafka.....	903
5.2 From MySQL to CSS/ES.....	918
5.3 From Oracle to Kafka.....	935
5.4 From Oracle to GaussDB Primary/Standby.....	951
5.5 From Oracle to GaussDB Distributed.....	974
5.6 From PostgreSQL to Kafka.....	997
5.7 From GaussDB Primary/Standby to Oracle.....	1010
5.8 From GaussDB Primary/Standby to Kafka.....	1029
5.9 From GaussDB Distributed to Oracle.....	1043
5.10 From GaussDB Distributed to Kafka.....	1063
6 Task Management.....	1078
6.1 Creating a Synchronization Task.....	1078
6.2 Querying the Synchronization Progress.....	1093
6.3 Viewing Synchronization Logs.....	1095

6.4 Comparing Synchronization Items.....	1096
6.5 Managing Objects.....	1108
6.5.1 Editing Synchronization Objects.....	1108
6.5.2 Importing Synchronization Objects.....	1110
6.5.3 Mapping Object Names.....	1111
6.5.4 Viewing Synchronization Mapping Information.....	1116
6.5.5 Processing Data.....	1117
6.6 Task Life Cycle.....	1126
6.6.1 Viewing Task Details.....	1126
6.6.2 Modifying Task Information.....	1126
6.6.3 Modifying Connection Information.....	1128
6.6.4 Modifying the Flow Control Mode.....	1129
6.6.5 Editing a Synchronization Task.....	1133
6.6.6 Resuming a Synchronization Task.....	1134
6.6.7 Skipping Data and Resuming a Synchronization Task.....	1135
6.6.8 Pausing a Synchronization Task.....	1135
6.6.9 Resetting a Synchronization Task.....	1137
6.6.10 Skipping a Synchronization Task.....	1140
6.6.11 Performing a Primary/Standby Switchover.....	1140
6.6.12 Cloning a Synchronization Task.....	1142
6.6.13 Changing Specifications.....	1143
6.6.14 Unsubscribing from a Yearly/Monthly Task.....	1144
6.6.15 Stopping a Synchronization Task.....	1146
6.6.16 Deleting a Synchronization Task.....	1147
6.6.17 Task Statuses.....	1148
7 Tag Management.....	1151
8 Diagnosis of Abnormal Records.....	1154
9 Interconnecting with CTS.....	1155
9.1 Key Operations Recorded by CTS.....	1155
9.2 Viewing Traces.....	1155
10 Interconnecting with Cloud Eye.....	1157
10.1 Supported Metrics.....	1157
10.2 Configuring Alarm Rules.....	1162
10.3 Viewing Monitoring Metrics.....	1163
11 Operation Reference in Synchronization Scenarios.....	1165
11.1 Kafka Message Format.....	1165
11.2 Kafka Authentication.....	1171
11.3 Forcibly Stopping Synchronization of GaussDB Distributed.....	1174
11.4 Forcibly Stopping Synchronization of GaussDB Primary/Standby.....	1175
11.5 Forcibly Stopping Synchronization of PostgreSQL.....	1176

11.6 Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL 1178

A Change History..... 1181

1 Synchronization Overview

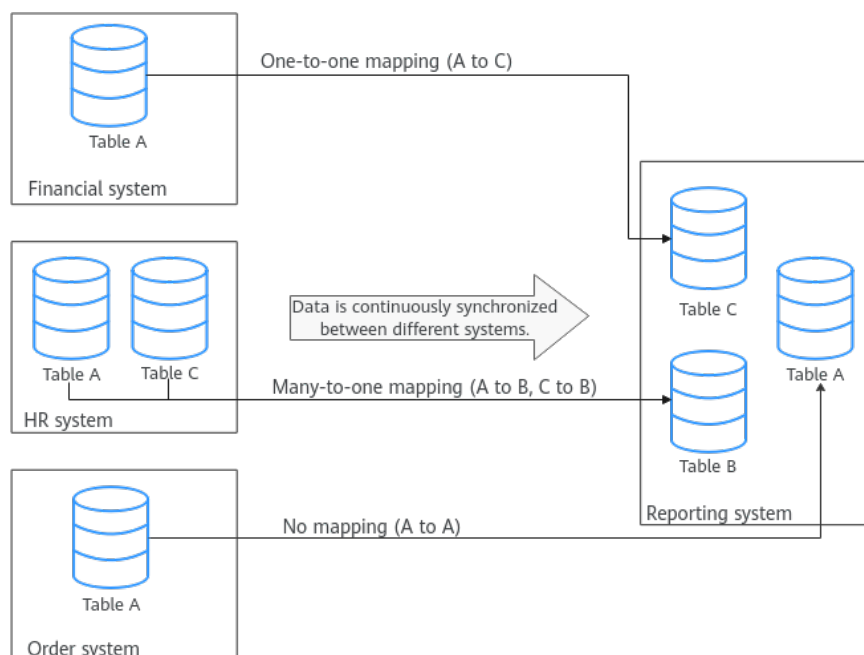
Real-time synchronization refers to the real-time flow of key service data from sources to destinations through a synchronization instance while consistency of data can be ensured.

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 Many-to-one real-time synchronization process



Supported Database Types

Table 1-1 lists the source database and destination database types supported by DRS in real-time migration.

Table 1-1 Synchronization scheme

Source DB	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	Incremental Full+Incremental	From MySQL to MySQL (To the cloud)
	GaussDB(for MySQL)	Incremental Full+Incremental NOTE Only whitelisted users can perform the incremental synchronization.	From MySQL to GaussDB(for MySQL)
	RDS for PostgreSQL	Full Full+Incremental	From MySQL to PostgreSQL
	GaussDB distributed	Full+Incremental	From MySQL to GaussDB Distributed
	GaussDB primary/standby NOTE Only whitelisted users can use this function.	Full+Incremental	From MySQL to GaussDB Primary/Standby
	GaussDB(DWS)	Incremental Full Full+Incremental NOTE Only whitelisted users can perform the incremental synchronization.	From MySQL to GaussDB(DWS)
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)

Source DB	Destination DB Type	Synchronization Mode	Related Documents
	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 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)
<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)	Full+Incremental	From PostgreSQL to GaussDB(DWS)
	GaussDB primary/standby	Incremental Full Full+Incremental NOTE Only whitelisted users can perform the incremental synchronization.	From PostgreSQL to GaussDB Primary/Standby
	GaussDB distributed	Incremental Full Full+Incremental NOTE Only whitelisted users can perform the incremental synchronization.	From PostgreSQL to GaussDB Distributed

Source DB	Destination DB Type	Synchronization Mode	Related Documents
<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)
<ul style="list-style-type: none"> RDS for PostgreSQL 	Kafka	Incremental	From PostgreSQL to Kafka (Out of the cloud)
DDM	RDS for MySQL	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 	Full+Incremental	From DDM to MySQL (Out of the cloud)
	DDM	Full+Incremental	From DDM to DDM
	GaussDB(DWS)	Full+Incremental	From DDM to GaussDB(DWS)
	<ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS 	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	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	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)	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 	Incremental	From GaussDB Distributed to MySQL
	<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)

Source DB	Destination DB Type	Synchronization Mode	Related Documents
	GaussDB(DWS)	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
	GaussDB primary/standby NOTE Only whitelisted users can use this function.	Full Incremental Full+Incremental	From GaussDB Distributed to GaussDB Primary/Standby
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 NOTE Only whitelisted users can use this function.	Full Incremental Full+Incremental	From GaussDB Primary/Standby to MySQL
	<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)

Source DB	Destination DB Type	Synchronization Mode	Related Documents
	Kafka	Incremental	From GaussDB Primary/Standby to Kafka (Out of the cloud) From GaussDB Primary/Standby to Kafka (self-built - self-built)
	GaussDB(DWS)	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB(DWS)
	GaussDB distributed	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB Distributed
	GaussDB primary/standby	Full Incremental Full+Incremental	From GaussDB Primary/Standby to GaussDB Primary/Standby
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)	Full+Incremental	From GaussDB(for MySQL) to GaussDB(DWS)
	Kafka	Incremental Full+Incremental	From GaussDB(for MySQL) to Kafka
	CSS/ES	Full+Incremental	From GaussDB(for MySQL) to CSS/ES

Source DB	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
DB2 for LUW	GaussDB distributed NOTE Only whitelisted users can use this function.	Full Full+Incremental	From DB2 for LUW to GaussDB Primary/Standby
	GaussDB primary/standby NOTE Only whitelisted users can use this function.	Full Full+Incremental	From DB2 for LUW to GaussDB Distributed
	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full Full+Incremental	From DB2 for LUW to GaussDB(DWS)
TiDB	GaussDB(for MySQL) NOTE Only whitelisted users can use this function.	Full+Incremental	From TiDB to GaussDB(for MySQL)
<ul style="list-style-type: none"> On-premises Microsoft SQL Server database 	GaussDB(DWS) NOTE Only whitelisted users can use this function.	Full+Incremental	From Microsoft SQL Server to GaussDB(DWS)

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


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. Multiple tables can be synchronized to one table.

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 databases on an ECS• MySQL databases on other clouds• RDS for MySQL	<ul style="list-style-type: none">• RDS for MySQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 on source databases. 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 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.

Table 3-2 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: 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.

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, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects. ● 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. ● Tables with storage engine different to MyISAM and InnoDB 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 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.

Type	Restrictions
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 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. ● 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.

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 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. ● 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. ● 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 ● Only MySQL to MySQL synchronization supports many-to-one synchronization. During table-level many-to-one synchronization, tables without primary keys cannot exist in the source database. ● If you create many-to-one synchronization tasks, the system automatically creates a parent task to associate multiple synchronization tasks after the tasks are started. The parent task is named in the <i>DRS-Group-Destination DB instance name</i> format. ● 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. ● 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 and destination sides are RDS for MySQL instances, tables encrypted using TDE cannot be synchronized. ● 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. ● 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.

Type	Restrictions
	<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. ● 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. ● 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: <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. ● During data synchronization, if the source database is on an RDS instance that does not belong the current cloud platform, the IP address cannot be changed. If the source database is on an RDS DB instance on the current cloud platform, the system automatically changes the IP address to the correct one. Then, retry the task to continue the synchronization. Therefore, changing the IP address is not recommended. ● To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> DDL statements can be synchronized during incremental synchronization. Incremental synchronization supports table renaming. The source and destination tables must be selected in the synchronization object list. Full synchronization does not support the RENAME command. 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 objects during an incremental synchronization.

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

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields:

- Region:** A dropdown menu with a downward arrow.
- Project:** A dropdown menu with a downward arrow.
- * Task Name:** A text input field containing "DRS-6131" and a help icon.
- Description:** A large text area with a help icon and a character count "0/256" at the bottom right.

Table 3-3 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-2 Synchronization instance information

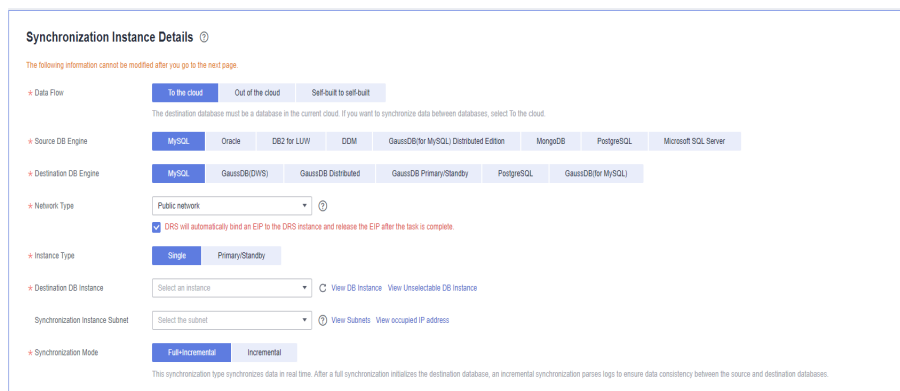


Table 3-4 Synchronization instance settings

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	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect

Parameter	Description
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none"> - Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. - Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</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 Subnet 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 ensure that there are available IP addresses. 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. <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> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-3 AZ</p>

- DRS instance specifications

Figure 3-4 Specifications



Table 3-5 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-5 Enterprise projects and tags

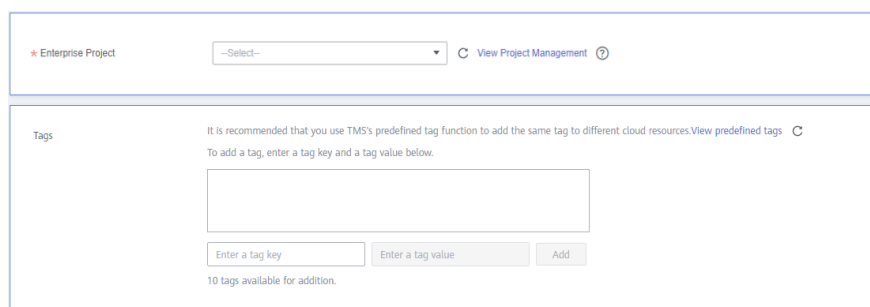


Table 3-6 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

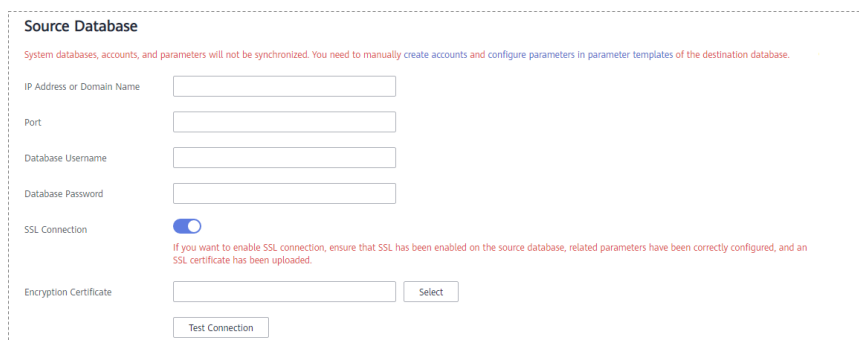
If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

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 3-6 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.

IP Address or Domain Name

Port

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 3-7 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. 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	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 not enabled, 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 3-7 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection


 Test successful

Table 3-8 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 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	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 not enabled, 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 3-8 Synchronization mode

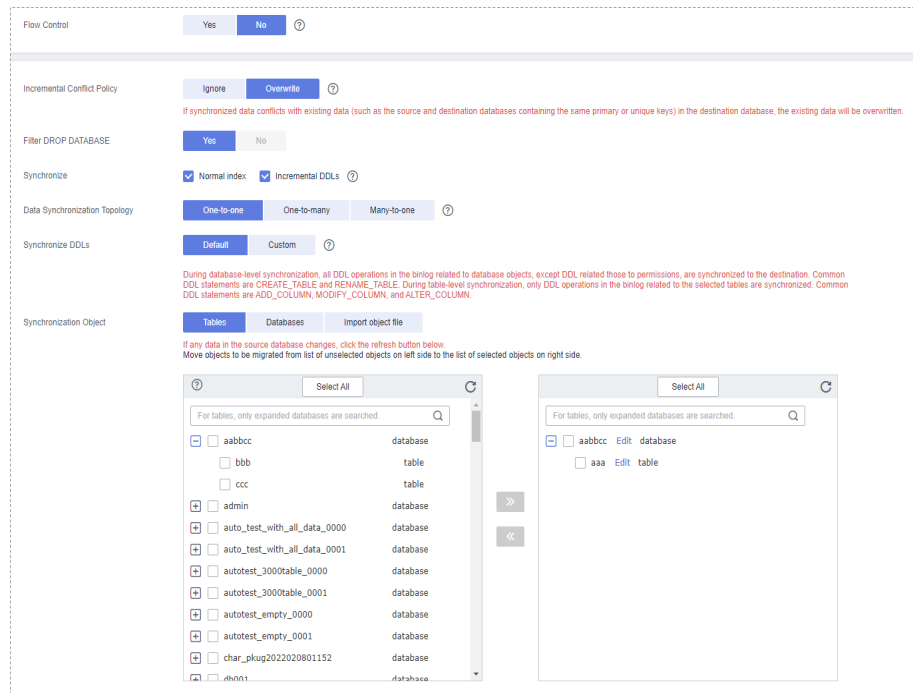
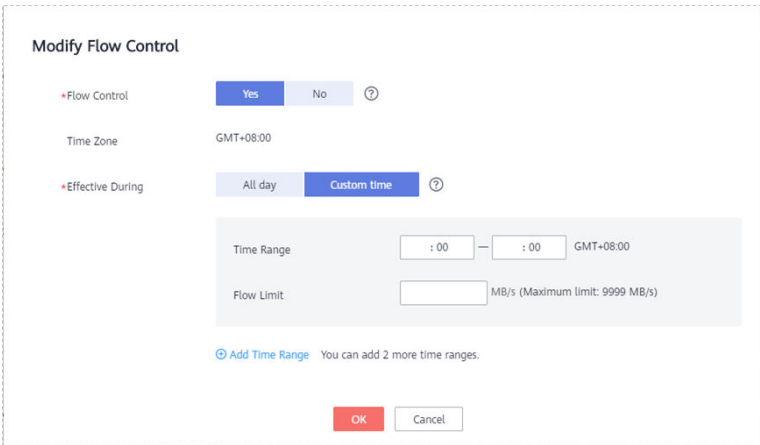



Table 3-9 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-9 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
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 data based on service requirements.</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 start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</p>

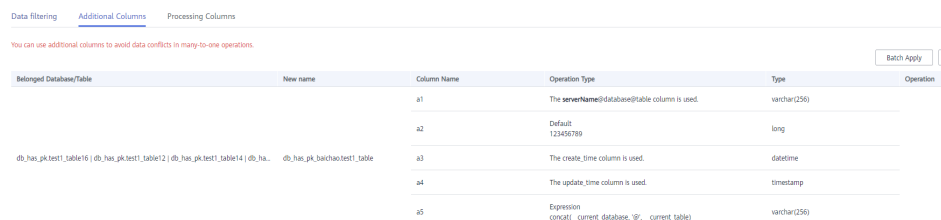
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: 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. <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 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 Column**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-10 Processing data




Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
		a1	The serverName@database@table column is used.	varchar(256)	
		a2	Default 123456789	long	
db_has_pk.test1_table16 db_has_pk.test1_table12 db_has_pk.test1_table14 db_has_pk.test1_table18		a3	The create_time column is used.	datetime	
		a4	The update_time column is used.	timestamp	
		a5	Expression concat(_current_database, '_')	varchar(256)	

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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 3-11 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-12 Task startup settings

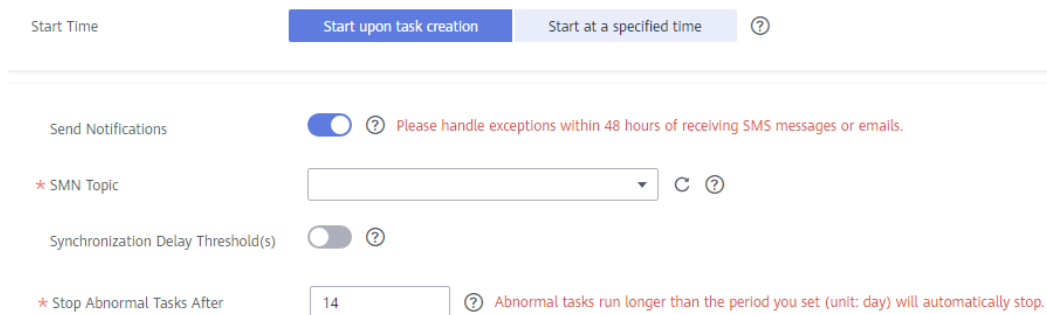



Table 3-10 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.2 From MySQL to PostgreSQL

Supported Source and Destination Databases

Table 3-11 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises MySQL databases• MySQL databases on an ECS• MySQL databases on other clouds• RDS for MySQL	<ul style="list-style-type: none">• RDS for PostgreSQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 3-12 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> • The source database user must have the following minimum permissions. <ul style="list-style-type: none"> - Minimum permission for full synchronization: SELECT. - Minimum permission for full and incremental synchronization: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the minimum permissions: the default permissions of an RDS PostgreSQL instance account

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> • Only tables, indexes, and (primary key, unique, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, and virtual columns cannot be synchronized. • MySQL views support syntax "as select ... from a join b where ..." but PostgreSQL does not, which may lead to the synchronization task to fail. • The following data types are not supported: XML, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and json. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. • 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.
Source database	<ul style="list-style-type: none"> • The names of the source databases and tables cannot contain non-ASCII characters, or special characters <'>.`\`" • 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. • The server_id value of the source MySQL database must be in the range from 1 to 4294967296.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly. If the destination DB instance type is RDS primary/standby instance, the replication status must also be normal.• 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.

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 associations: tables or views referenced by views ● If the source and destination sides are RDS for MySQL instances, tables encrypted using TDE cannot be synchronized. ● 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. ● 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. ● You can map multiple databases to one database for synchronization. Tables with the same name cannot exist in mapped databases. To use this function, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● 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. ● Different types of indexes synchronized to the destination database will become B-Tree indexes. ● 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. ● 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.

Type	Restrictions
	<ul style="list-style-type: none"> ● 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 source database cannot be restored. ● 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. ● 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 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. ● DDL operations are not supported during full synchronization. ● During incremental synchronization, 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. - Table fields can be added, deleted, and modified. For example: <pre style="background-color: #f0f0f0; padding: 5px;">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 style="background-color: #f0f0f0; padding: 5px;">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. If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. For example:

Type	Restrictions
	<pre>alter table `ddl_test` modify column `c1` varchar(50);</pre> <ul style="list-style-type: none"> - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. • 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. • 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. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</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 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 3-13 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is an orange warning banner with the text: "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." Below the banner, the form contains the following fields:

- Region:** A dropdown menu with a question mark icon.
- Project:** A dropdown menu with a question mark icon.
- Task Name:** A text input field containing "DRS-8131" and a question mark icon.
- Description:** A text area with a question mark icon and a character count "0/256" at the bottom right.

Table 3-13 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-14 Synchronization instance details

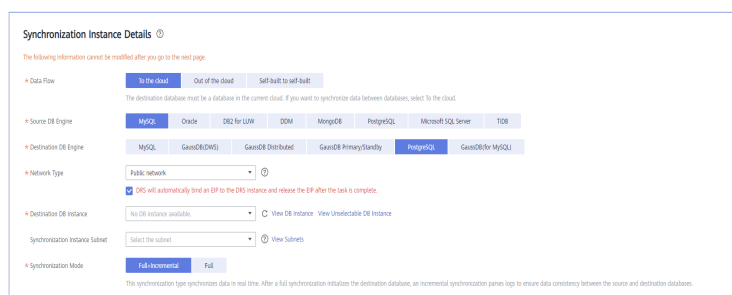


Table 3-14 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select PostgreSQL .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example.
Destination DB Instance	The RDS PostgreSQL DB instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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. - 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.

- DRS instance specifications

Figure 3-15 Specifications



Table 3-15 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-16 Enterprise projects and tags




Table 3-16 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-17 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 3-17 Self-build on ECS - 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.
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 not enabled, 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 3-18 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-18 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 3-19 Synchronization mode

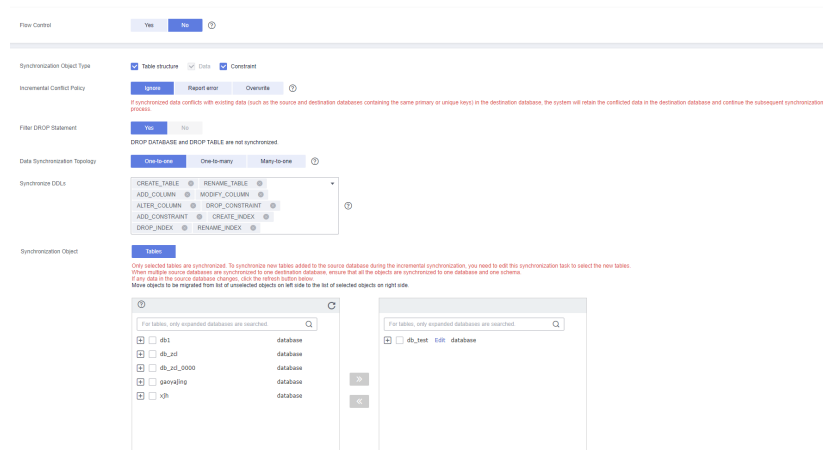
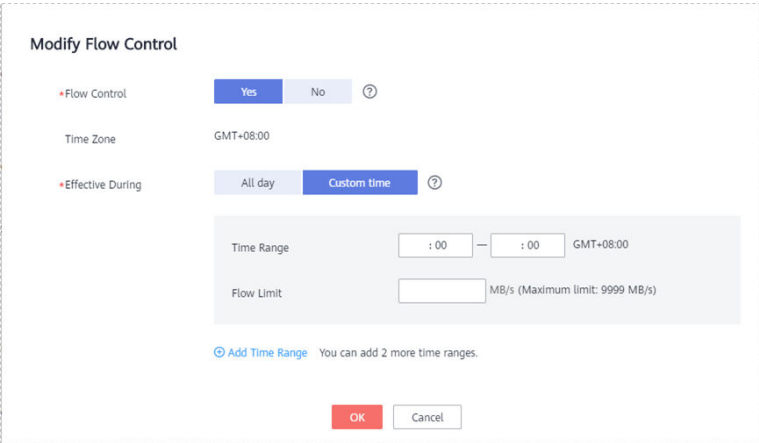



Table 3-19 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-20 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>Available options: Table structure, Data, and Constraint</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.
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.
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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</p>

Parameter	Description
Synchronize DDLs	<p>You can select the DDL types for incremental synchronization as required. The DDL type depends on the application scenario.</p> <ul style="list-style-type: none"> 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. <p>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.</p>
Synchronize	Incremental DDL synchronization is supported. You can determine whether to synchronize data based on service requirements.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

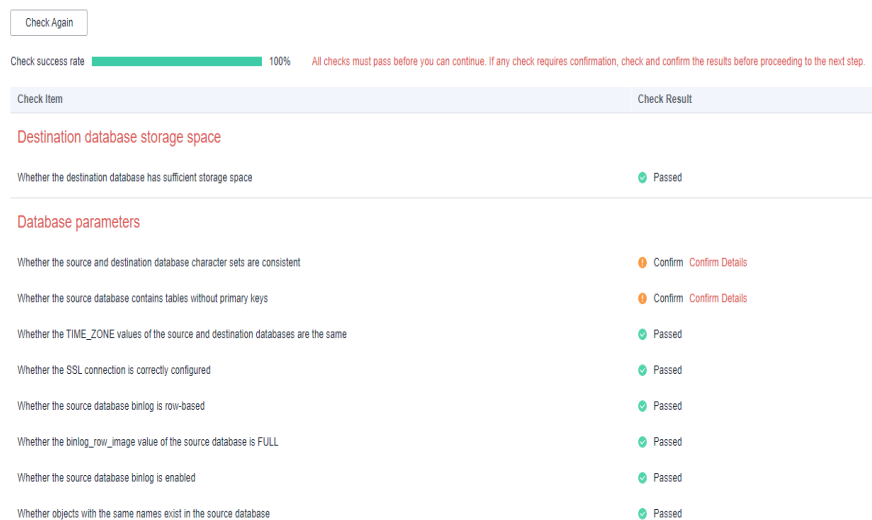
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-21 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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

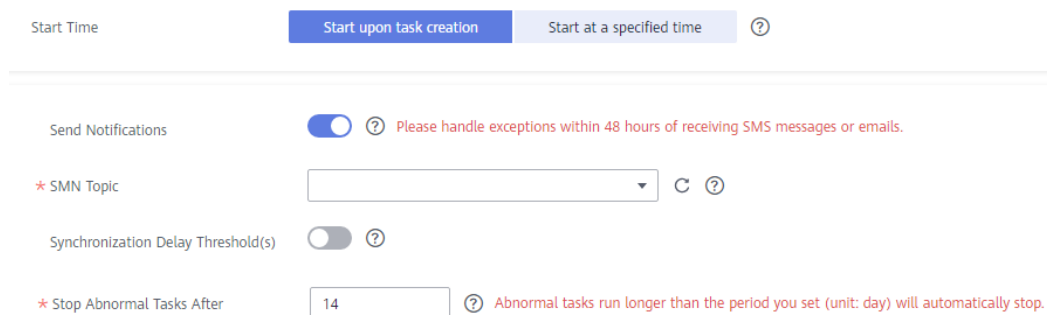



Table 3-20 Task startup settings

Parameter	Description
Started 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>SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.</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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.3 From MySQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-21 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECSMySQL databases on other cloudsRDS for MySQL	GaussDB distributed

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 3-22 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 GaussDB database user must have the following permissions: <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. - Schema-level permission: 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 permission: 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.
Synchronization object	<ul style="list-style-type: none"> ● Only tables can be synchronized. Other database objects such as stored procedures cannot be synchronized. ● Only tables with primary keys can be synchronized. Tables without primary keys cannot be synchronized. ● Incremental synchronization does not support synchronization of DDL. ● MySQL tables containing virtual columns cannot be synchronized. ● The following data types are not supported: XML, JSON, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● 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 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. For MySQL 5.7, the value of server_id ranges from 1 to 4294967296. ● The names of the source databases and tables cannot contain non-ASCII characters, or special characters .<'>\\
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 mapped database configured for the task must exist in the destination 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. ● 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. ● 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. ● 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. ● After a table in the source database is synchronized to the destination database, the table is distributed in hash mode and cannot be replicated. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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. ● 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 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. ● 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

Type	Restrictions
	<p>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.</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, DDLs of the source database cannot be replicated.• During an incremental synchronization, do not perform the restoration operation on the source database.• During task startup or full synchronization, you are not advised to perform DDL operations, such as the deletion operation. 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 cannot be restored.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• 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.• 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.• Two-phase commit is not supported.• The ongoing synchronization task cannot be paused.• Data processing is not supported.• 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 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

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning banner with a red triangle icon and the text: "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." Below the banner, there are four input fields:

- Region:** A dropdown menu with a downward arrow.
- Project:** A dropdown menu with a downward arrow.
- + Task Name:** A text input field containing "DRS-6131" and a help icon.
- Description:** A text area with a height of 0/256 characters and a help icon.

Table 3-23 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-24 Synchronization instance details

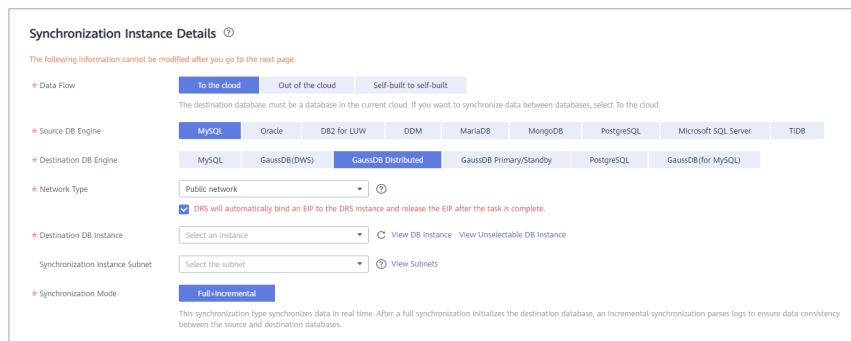


Table 3-24 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
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
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 Subnet 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 ensure that there are available IP addresses. 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. 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.

- Enterprise Projects and Tags

Figure 3-25 Enterprise projects and tags




Table 3-25 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

- 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 3-26 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 3-26 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 not enabled, 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 3-27 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields: "DB Instance Name", "Database Username", and "Database Password". Below these fields is a "Test Connection" button. The form is enclosed in a dashed border.

Table 3-27 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 3-28 Synchronization mode

The screenshot displays the "Synchronization mode" configuration page. It includes three main sections:

- Incremental Conflict Policy:** Radio buttons for "Ignore", "Report error", and "Overwrite". A tooltip explains: "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:** Radio buttons for "Yes" and "No".
- Synchronization Object:** A radio button for "Tables". A tooltip states: "This task does not support DDL synchronization. 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."

 At the bottom, there are two panels for object selection. The left panel shows a list of objects: "db_zkn1_0000" (database), "mycat_sequence" (table), and "tableaddzknpart2" (table). The right panel shows a search bar and a list of selected objects, including "db_zkn1_0000" with an "Edit" link. Navigation arrows are present between the panels.

Table 3-28 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. • 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.
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select table-level 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

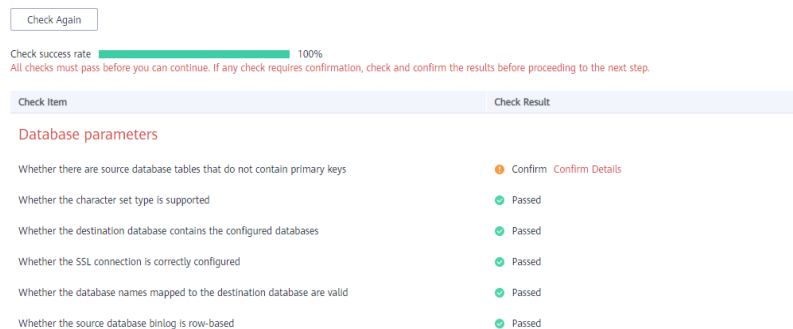
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-29 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-30 Task startup settings

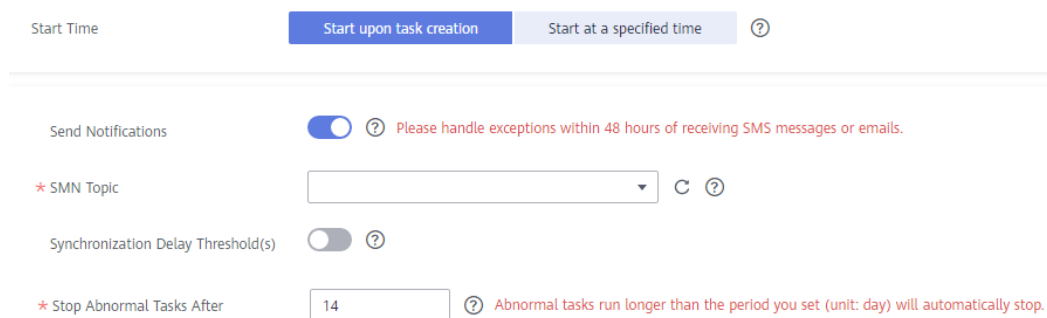



Table 3-29 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.4 From MySQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-30 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises MySQL databases• MySQL databases on an ECS• MySQL databases on other clouds• RDS for MySQL	GaussDB primary/standby

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 3-31 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 GaussDB database user must have the following permissions: <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. - Schema-level permission: 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 permission: 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.
Synchronization object	<ul style="list-style-type: none"> • Only tables can be synchronized. Other database objects such as stored procedures cannot be synchronized. • Incremental synchronization does not support synchronization of DDL. • MySQL tables containing virtual columns cannot be synchronized. • The following data types are not supported: XML, JSON, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● 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 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. For MySQL 5.7, the value of server_id ranges from 1 to 4294967296. ● The names of the source databases and tables cannot contain non-ASCII characters, or special characters .<'>\/
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 mapped database configured for the task must exist in the destination 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. ● 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. ● 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. ● 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. ● After a table in the source database is synchronized to the destination database, the table is distributed in hash mode and cannot be replicated. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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. ● 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 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. ● 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

Type	Restrictions
	<p>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.</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, DDLs of the source database cannot be replicated. ● During an incremental synchronization, do not perform the restoration operation on the source database. ● During task startup or full synchronization, you are not advised to perform DDL operations, such as the deletion operation. 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 cannot be restored. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● 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. ● 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. ● Two-phase commit is not supported. ● Data processing is not supported. ● 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 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 3-31 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "Task Name" (a text input field containing "DRS-6131"), and "Description" (a text area with a character count of "0/256"). A small note below the Region field states: "Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region."

Table 3-32 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-32 Synchronization instance details

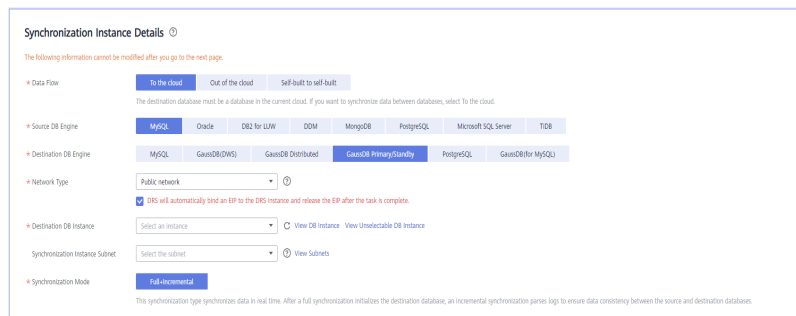


Table 3-33 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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. 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.

- Enterprise Projects and Tags

Figure 3-33 Enterprise projects and tags

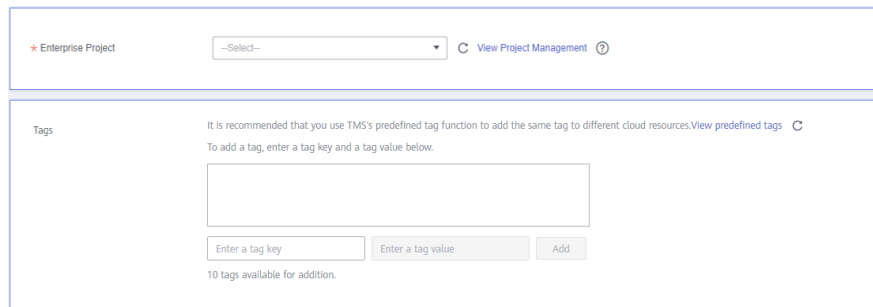


Table 3-34 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-34 Source database information




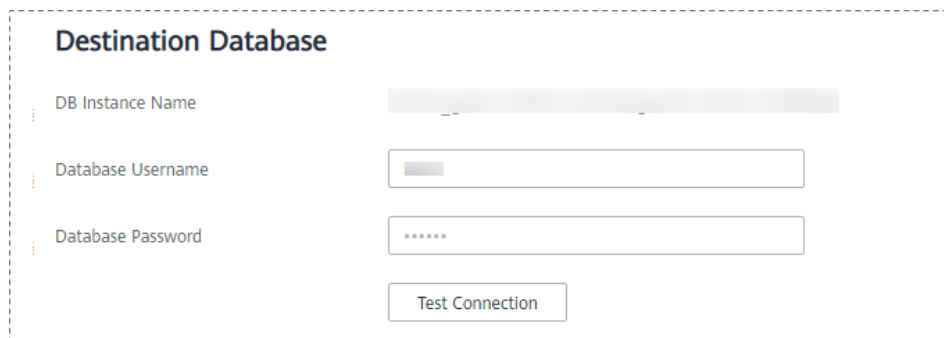
Table 3-35 Source database settings

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 not enabled, 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 3-35 Destination database information



The screenshot shows a configuration form titled "Destination Database". It contains three input fields: "DB Instance Name" (with a greyed-out value), "Database Username" (with a greyed-out value), and "Database Password" (with masked characters). Below these fields is a "Test Connection" button.

Table 3-36 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.

Parameter	Description
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 3-36 Synchronization mode

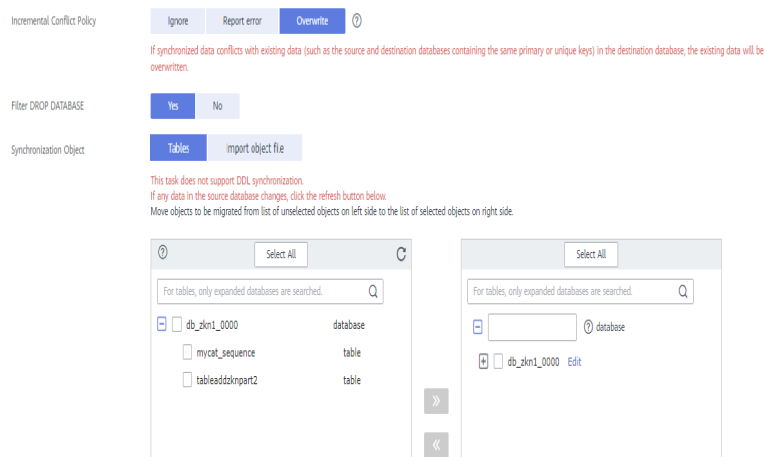
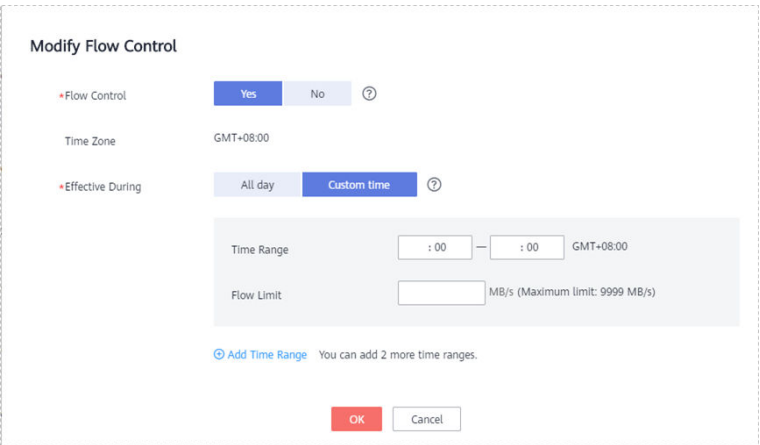



Table 3-37 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-37 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

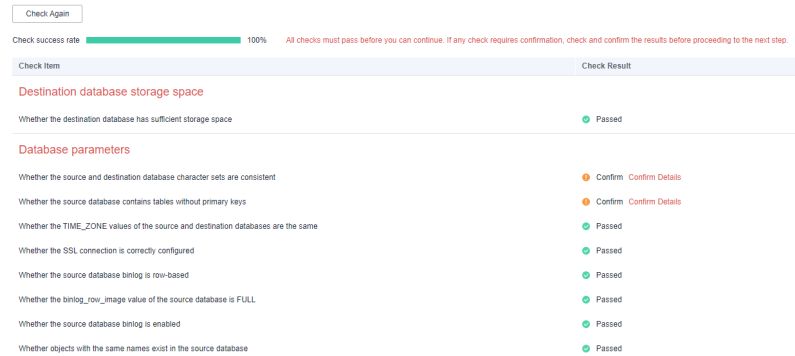
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-38 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-39 Task startup settings

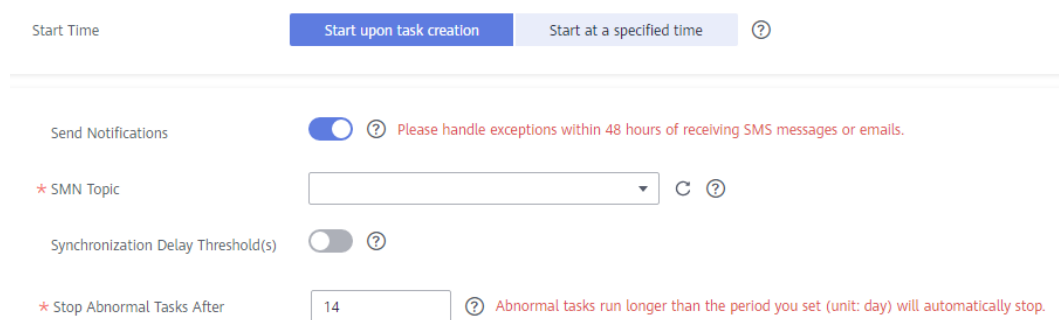



Table 3-38 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.5 From MySQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-39 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECSMySQL databases on other cloudsRDS for MySQL	<ul style="list-style-type: none">GaussDB(DWS) cluster 8.1.3

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 3-40 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, CREATE, and REFERENCES.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, indexes, and constraints (primary key, null, not null) can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraint, and unique index 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. ● Tables with storage engine different to MyISAM and InnoDB cannot 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 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.
Source database	<ul style="list-style-type: none"> ● 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. ● The server_id value of the source MySQL database must be in the range from 1 to 4294967296. ● The names of the source databases and tables cannot contain non-ASCII characters, or special characters <'>.\ "

Type	Restrictions
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.

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. ● 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. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● 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 DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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). ● 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. ● 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. ● 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.

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. ● 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. ● 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. ● During incremental synchronization, the following 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, ALTER_COLUMN, CREATE_INDEX, DROP_INDEX, RENAME_INDEX, CHANGE_COLUMN, TRUNCATE_TABLE and DROP_TABLE. 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. – The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails.

Type	Restrictions
	<ul style="list-style-type: none"> - 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. • 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, tables without primary keys created during database-level synchronization must contain one of the following. Integer types: TINYINT, SMALLINT, INT, BIGINT, and NUMERIC/DECIMAL Character types: CHAR, BPCHAR, VARCHAR, VARCHAR2, and NVARCHAR2 Date/time types: DATE, TIME, TIMETZ, TIMESTAMP, TIMESTAMPZ, 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"> • 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. • 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 the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the GaussDB(DWS) 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 3-40 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-6131 [Help icon]

Description: [Text area, 0/256]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Table 3-41 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-41 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 | MongoDB | PostgreSQL | Microsoft SQL Server | TiDB

* Destination DB Engine: MySQL | **GaussDB(DWS)** | GaussDB Distributed | GaussDB Primary/Standby | PostgreSQL | 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: VPC [Help icon]

* Instance Type: **Single** | Primary/Standby

* Destination DB Instance: [Dropdown menu] | View DB Instance | View Unselectable DB Instance

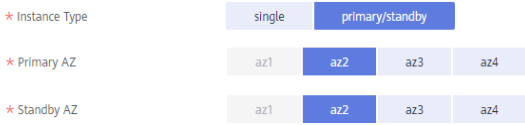
Synchronization Instance Subnet: [Dropdown menu] | View Subnets

* Synchronization Mode: **Full-Incremental** | 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.

Table 3-42 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	The VPC network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none">- Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization.- Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</p>
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 Subnet 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 ensure that there are available IP addresses. 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. In this section, 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 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.
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-42 AZ</p>  <p>The screenshot shows a configuration interface with three rows of options:</p> <ul style="list-style-type: none"> Instance Type: single (grey), primary/standby (blue) Primary AZ: az1 (grey), az2 (blue), az3 (grey), az4 (grey) Standby AZ: az1 (grey), az2 (blue), az3 (grey), az4 (grey)

- DRS instance specifications

Figure 3-43 Specifications



Table 3-43 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-44 Enterprise projects and tags




Table 3-44 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 3-45 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.

Source Database Type: Self-built on ECS RDS DB Instance

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

Database Username:

Database Password:

Table 3-45 Source database settings

Parameter	Description
Source Database Type	Select Self-built on ECS 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 3-46 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

Table 3-46 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 3-47 Synchronization mode

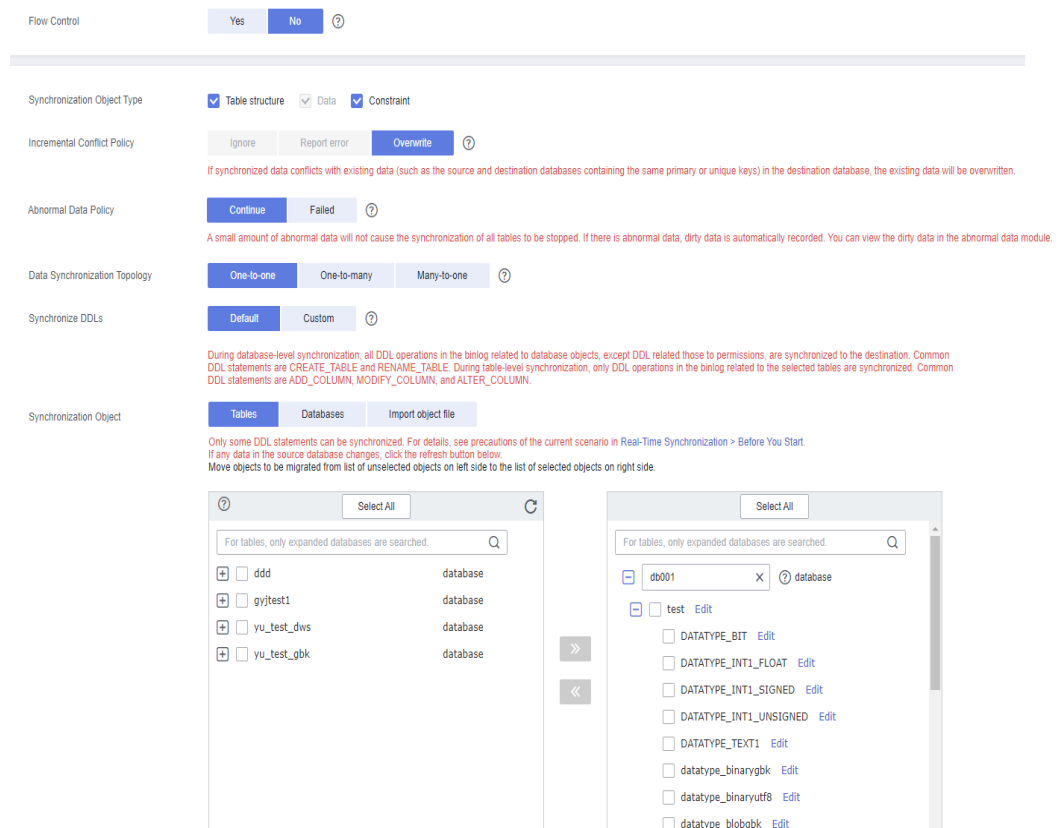
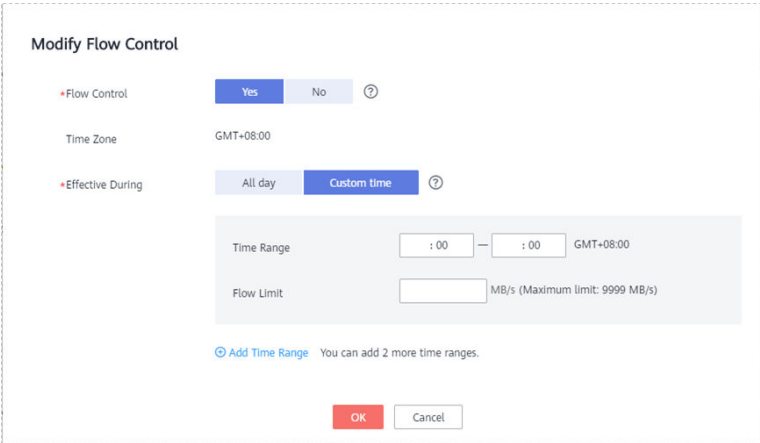



Table 3-47 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-48 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>Available options: Table structure, Data, and Constraint</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. • 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.
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</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. <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 to enable many-to-one synchronization. For details, see 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 3-49 Processing data

Original Database/Table	New name	Column Name	Operation Type	Type	Operation
db1_jack_011	db1_jack_011	--	--	--	Add
db1_jack2	db1_jack2	--	--	--	Add
db1_jack3	db1_jack3	--	--	--	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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-50 Pre-check

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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-51 Task startup settings

Start Time: Start upon task creation Start at a specified time

Send Notifications: Please handle exceptions within 48 hours of receiving SMS messages or emails.

SMN Topic:


Synchronization Delay Threshold(s):

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

Table 3-48 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.6 From MySQL to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-49 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises MySQL databases• MySQL databases on an ECS• MySQL databases on other clouds• RDS for MySQL	<ul style="list-style-type: none">• GaussDB(for MySQL)

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 3-50 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 GaussDB(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.
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, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects. • 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. • Tables with storage engine different to MyISAM and InnoDB 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 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.

Type	Restrictions
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: '<'>/'\ ● 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 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 or later, the server_id value ranges from 1 to 4294967296. ● 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.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly.• The destination DB instance must have sufficient storage 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. ● 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 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 and destination sides are RDS for MySQL instances, tables encrypted using TDE cannot be synchronized. ● 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. ● 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.

Type	Restrictions
	<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. • After a task is created, the destination database cannot be set to read-only. • 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. • 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.

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 3-52 Synchronization task information

The screenshot shows a web interface for creating a synchronization task. At the top, there is an orange warning banner with the text: "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." Below the banner, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "Task Name" (a text input field containing "DRS-B131"), and "Description" (a text area). A small note below the "Region" field reads: "Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region." At the bottom right of the form, there is a small "0/256" character count.

Table 3-51 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-53 Synchronization instance details

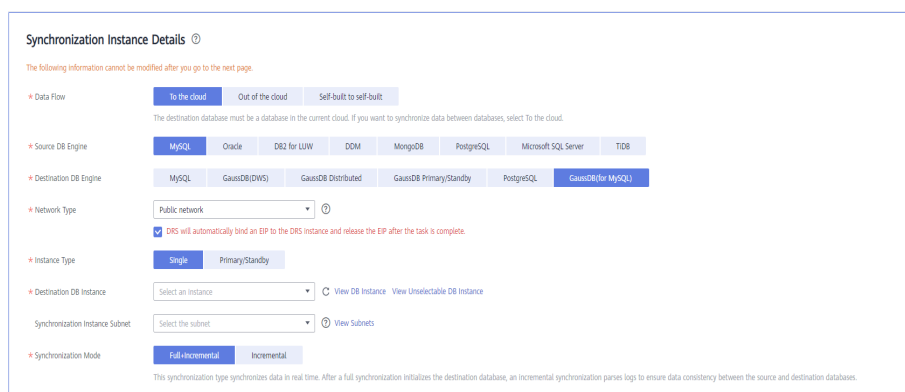
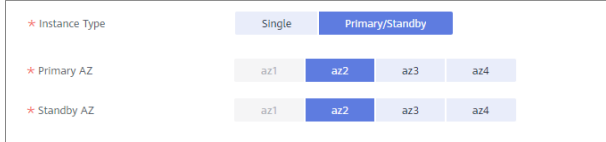


Table 3-52 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect

Parameter	Description
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none"> - Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. - Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</p>
Destination DB Instance	<p>An available GaussDB(for MySQL) instance.</p>
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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 <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>

Parameter	Description
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-54 AZ</p> 

- DRS instance specifications

Figure 3-55 Specifications



Table 3-53 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-56 Enterprise projects and tags




Table 3-54 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-57 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 3-55 Source database settings

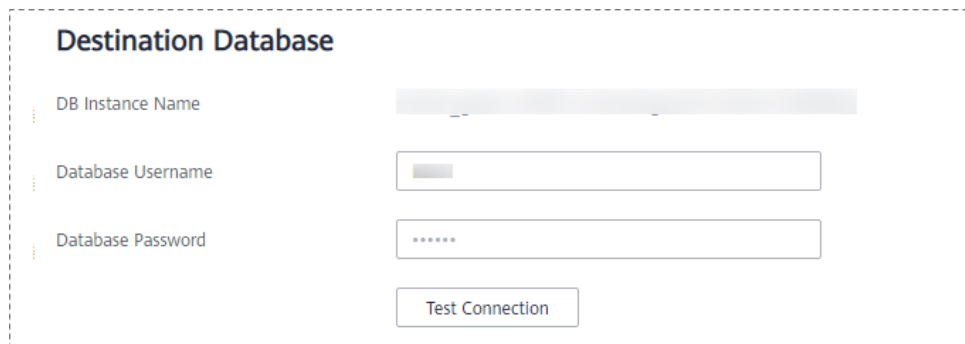
Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
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 not enabled, 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 3-58 Destination database information



The screenshot shows a form titled "Destination Database" with three input fields: "DB Instance Name", "Database Username", and "Database Password". Below these fields is a "Test Connection" button. The form is enclosed in a dashed border.

Table 3-56 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 3-59 Synchronization mode

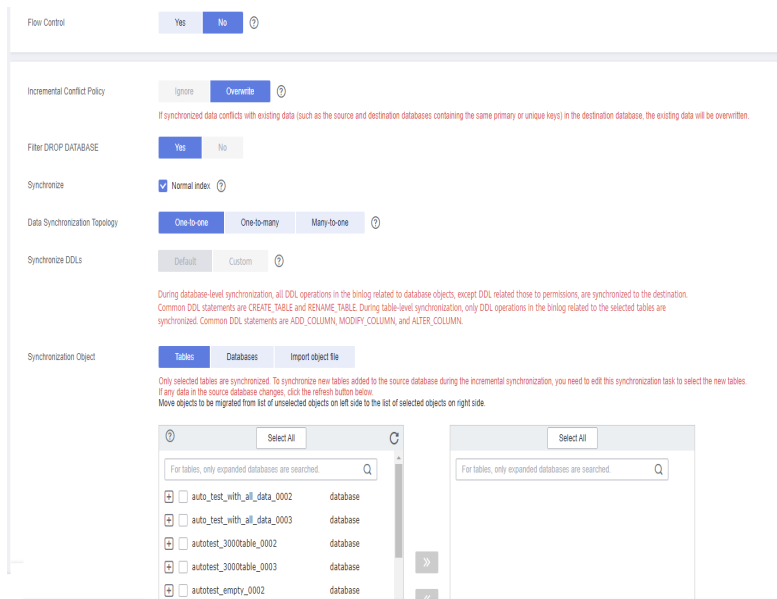
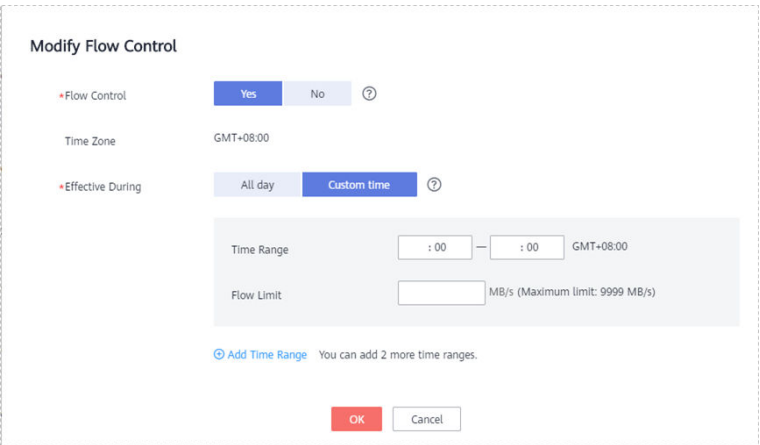



Table 3-57 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-60 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>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>
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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</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. <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.

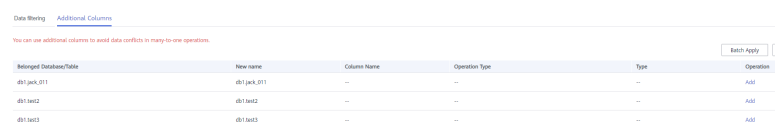
Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 3-61 Processing data



Original Database/Table	New name	Column Name	Operation Type	Type	Operation
db1_jack_011	db1_jack_011	--	--	--	Add
db1_said2	db1_said2	--	--	--	Add
db1_said3	db1_said3	--	--	--	Add

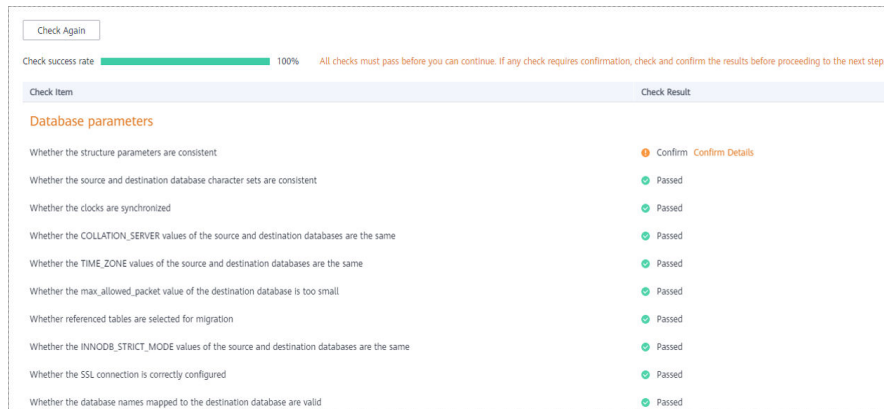
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-62 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-63 Task startup settings

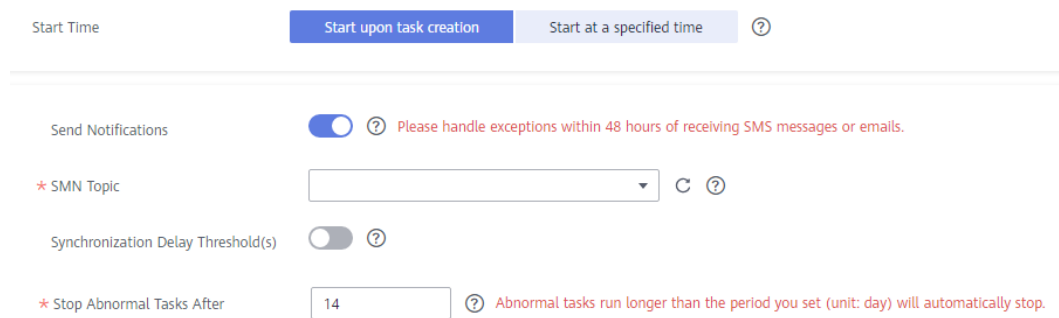



Table 3-58 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.7 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 3-59 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)	RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13, 14) NOTE The major version of the destination database must be the same as or later than that of the source database.

Supported Synchronization Objects

Table 3-60 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-60 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"> - Schema, table, index, constraint, view, materialized view, sequence, stored procedure, rule, trigger, foreign key, sorting rule, plug-in, code conversion information, aggregate function, operator, statistics extension, conversion information, text search configuration, function, data type, type conversion, user, and event trigger, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, and users 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

Type	Notes
	<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. • 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. • 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. • User: Existing users in the destination database, superuser, replication, and bypassrsls 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. <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 operations can be synchronized, including TRUNCATE (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, and ALTER VIEW. <p>During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (PostgreSQL 11 or later), DROP TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p>

Type	Notes
	<ul style="list-style-type: none">- Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE 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 Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</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 3-61](#). 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.

Table 3-61 Database account permission

Type	Full	Full+Incremental
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 md5</code> before all configurations in the pg_hba.conf file of the source database. 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full	Full+Incremental
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 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 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.</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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 3-62 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 style="background-color: #f0f0f0; padding: 2px;">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 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. ● 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 origin. ● 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,

Type	Constraints
	<p>including databases, schemas, and tables. System databases, system schemas, and system tables are excluded.</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: 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. <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. - 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. - If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database.

Type	Constraints
	<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 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.
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. • 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+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. – If the value of destination database session_replication_role is replica when the full+incremental synchronization task is complete, change the value to origin. ● 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 destination database session_replication_role is replica, change it to origin to forcibly stop the full+incremental synchronization task.
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.

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

This section uses 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 3-64 Synchronization task information

Table 3-63 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-65 Synchronization instance details

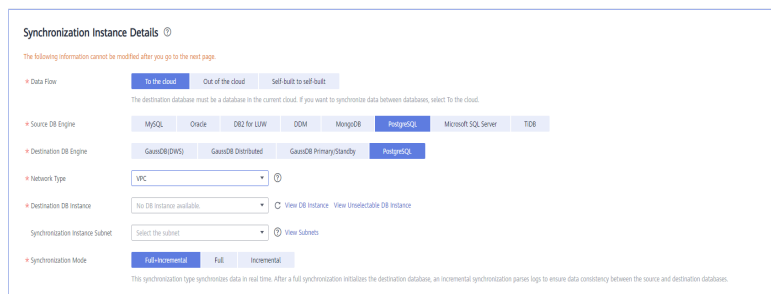


Table 3-64 Synchronization instance settings

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.
Destination DB Instance	The RDS for PostgreSQL DB instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 3-66 Specifications



Table 3-65 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-67 Enterprise projects and tags

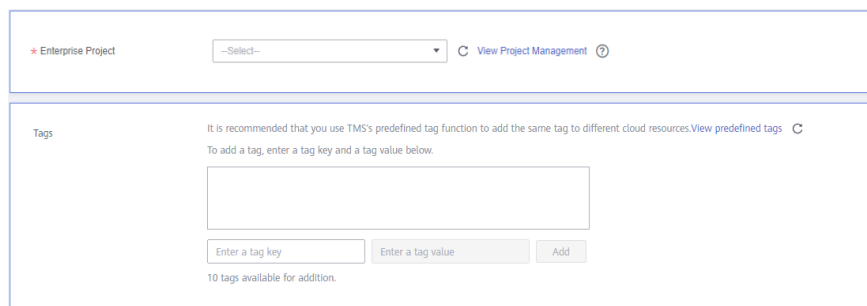


Table 3-66 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> – This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. – 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 ends.

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 3-68 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.

Source 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:

Table 3-67 Self-build 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.
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 3-69 RDS DB instance - source database information

The screenshot shows a configuration window titled "Source Database". At the top, there is a warning: "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." Below this, the "Source Database Type" is set to "RDS DB Instance". The "DB Instance Name" field has a dropdown menu and a "View DB Instance" link. The "Database Username" and "Database Password" fields are visible with masked characters. A "Test Connection" button is located at the bottom of the form.

Table 3-68 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 3-70 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields: "DB Instance Name", "Database Username", and "Database Password". Below the fields is a "Test Connection" button. The form is enclosed in a dashed border.

Table 3-69 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 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 3-71 Synchronization mode

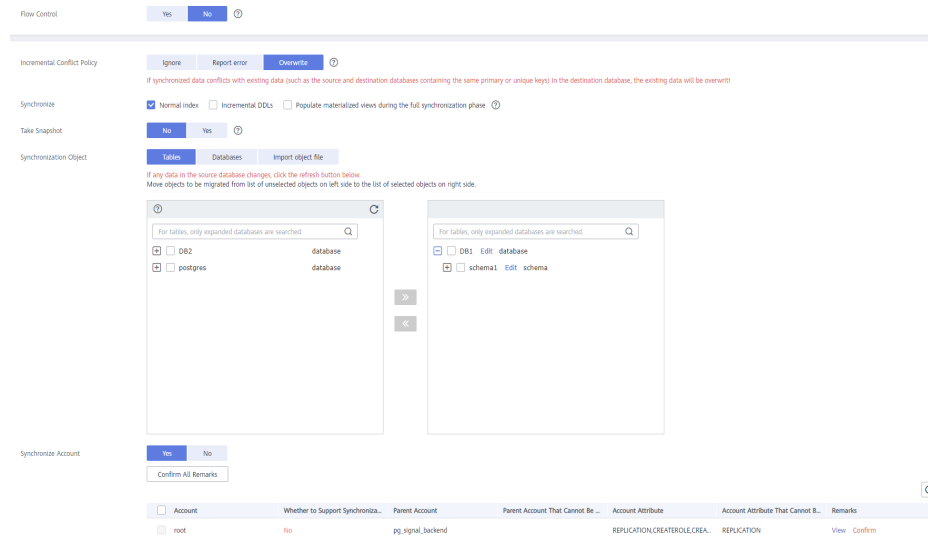
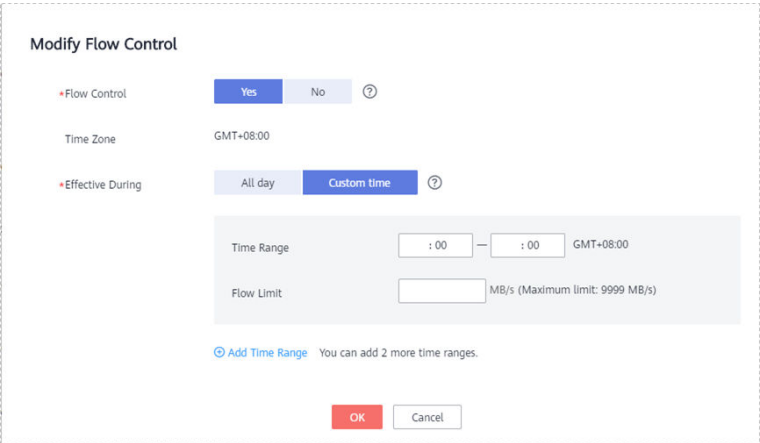



Table 3-70 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-72 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.
Synchronize	<p>Available options: 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>
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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.</p> <p>There are two types of accounts: accounts that can be synchronized and accounts that cannot be synchronized. For accounts that cannot be synchronized, you can view details in the Remarks column and determine whether to synchronize accounts and permissions based on your service requirements.</p>

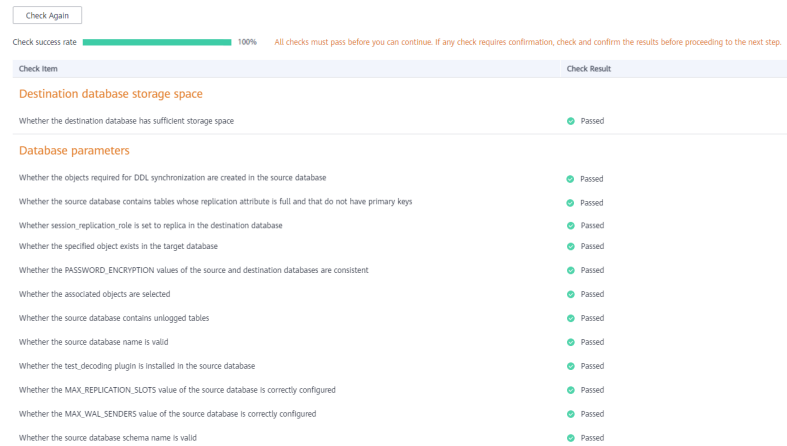
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-73 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-74 Task startup settings

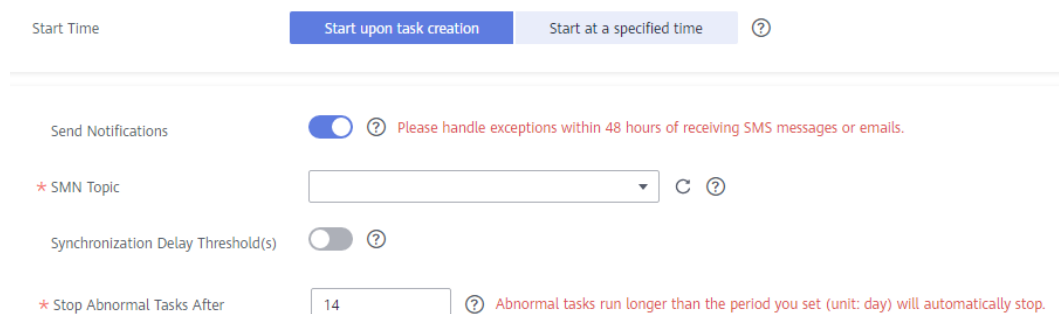



Table 3-71 Task startup settings

Parameter	Description
Started 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>SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.</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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.8 From PostgreSQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-72 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)	<ul style="list-style-type: none">• GaussDB(DWS) cluster

Supported Synchronization Objects

[Table 3-73](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-73 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"> - tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz, and interval - 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. <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. • 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

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 3-74](#). 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.

Table 3-74 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</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 md5</code> before all configurations in the pg_hba.conf file of the source database. 2. Run select pg_reload_conf(); 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 object 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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-75 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 style="background-color: #f0f0f0; padding: 2px;">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 associations are as follows: tables associated with primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by partitioned child tables, and sequences referenced by auto-increment columns</p> ● Destination database parameter requirements: <ul style="list-style-type: none"> – The block_size value of the destination database must be greater than that of the source database. ● 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 character set of the destination database must be the same as that of the source database.

Type	Constraints
	<ul style="list-style-type: none"> - 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: 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. <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. - 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 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 constraint 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	Constraints
	<ul style="list-style-type: none"> - 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.
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.
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. • 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.

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	<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.

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

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 3-75 Synchronization task information

Table 3-76 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-76 Synchronization instance details

Table 3-77 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . VPC is used as an example.
Destination DB Instance	The destination is a GaussDB(DWS) instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

- Enterprise Projects and Tags

Figure 3-77 Enterprise projects and tags




Table 3-78 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-78 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.

Source 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:

Table 3-79 Self-build 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.

Parameter	Description
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 3-79 RDS DB instance - 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.

Source Database Type: Self-built on ECS RDS DB Instance

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

Database Username:

Database Password:

✔ Test successful

Table 3-80 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 3-80 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

Table 3-81 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 3-81 Synchronization Mode

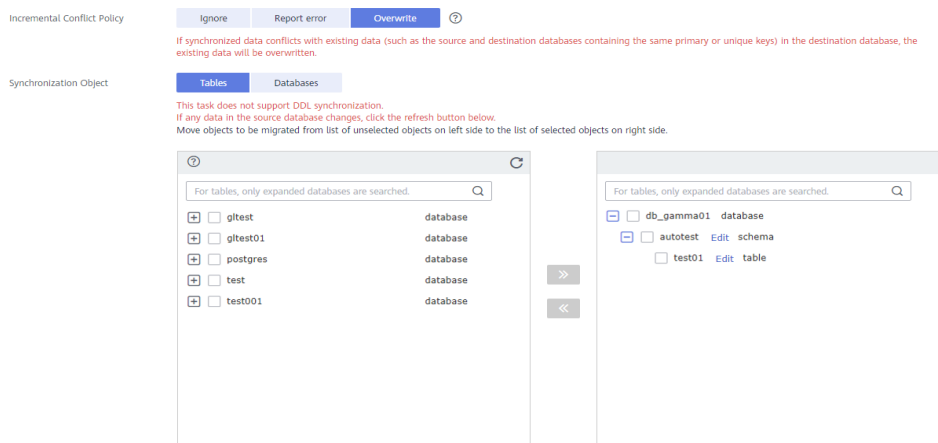



Table 3-82 Synchronization object

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. • Report error The synchronization task will be stopped and fail. • 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

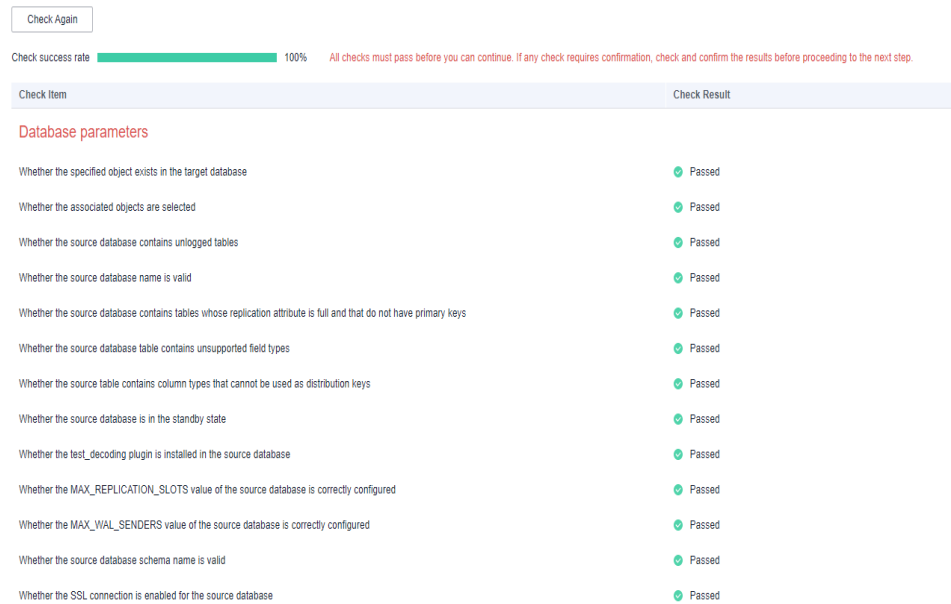
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-82 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-83 Task startup settings





Table 3-83 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.9 From PostgreSQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-84 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

Supported Synchronization Objects

Table 3-85 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-85 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"> - tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz, and interval - 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 keys 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"). <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

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 3-86](#). 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.

Table 3-86 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</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 md5</code> before all configurations in the pg_hba.conf file of the source database. 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Type	Full	Full+Incremental
Destination database user	<ul style="list-style-type: none"> ● 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. 	

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-87 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. - 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.
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.
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. ● 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.

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	<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.

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 3-84 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-6131

Description: [Text area]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Table 3-88 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-85 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 | MongoDB | PostgreSQL | Microsoft SQL Server | TiDB

* Destination DB Engine: GaussDB(DWS) | GaussDB Distributed | GaussDB Primary/Standby | PostgreSQL

* Network Type: VPC

* Destination DB Instance: Select an instance | View DB Instance | View Unselectable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnets

* 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.

Table 3-89 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
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.
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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing.

- Enterprise Projects and Tags

Figure 3-86 Enterprise projects and tags




Table 3-90 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-87 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.

Source 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:

Table 3-91 Self-build 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.

Parameter	Description
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 3-88 RDS DB instance - source database information

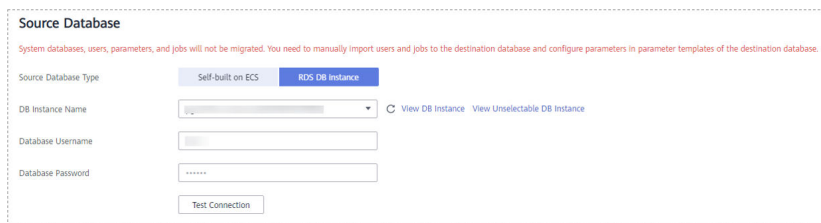


Table 3-92 RDS DB instance - source database information

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 3-89 Destination database information

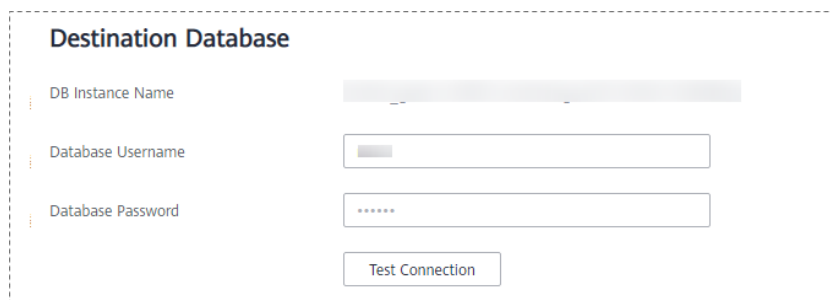


Table 3-93 Destination database settings

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 3-90 Synchronization mode

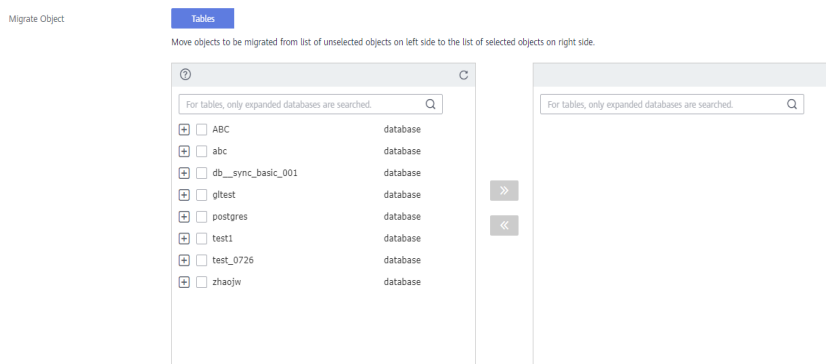



Table 3-94 Synchronization Object

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports 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 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. <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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

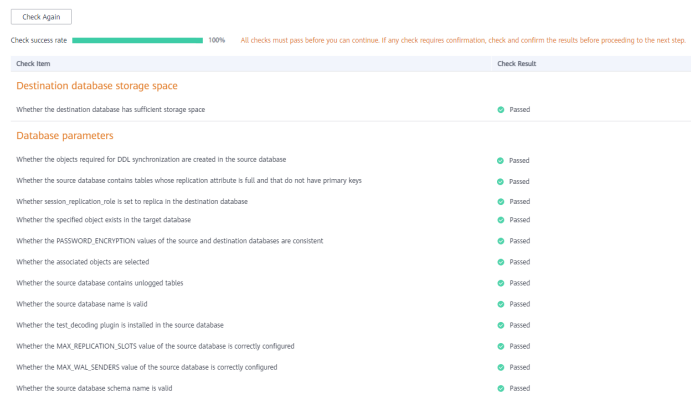
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-91 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-92 Task startup settings

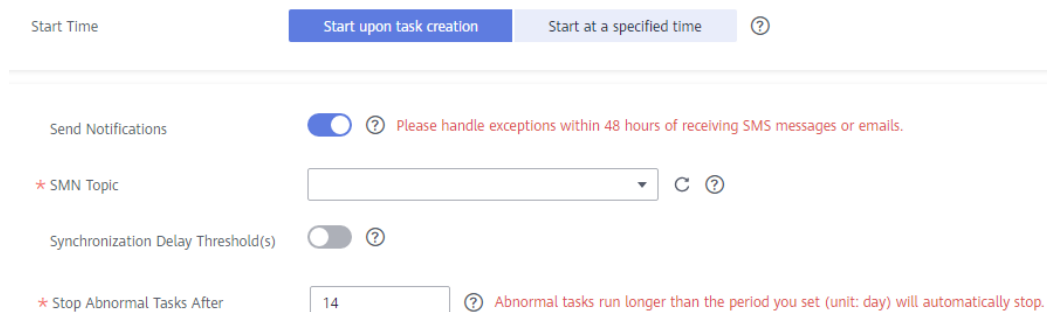



Table 3-95 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.10 From PostgreSQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-96 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

Supported Synchronization Objects

Table 3-97 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-97 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"> – tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz, and interval – 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 keys 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"). <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

Type	Precautions
	<ul style="list-style-type: none">- Not supported: DDL statements, DML statements of unlogged tables and temporary tables

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 3-98](#). 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.

Table 3-98 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</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> <drs_instance_ip>/32 md5</code> before all configurations in the pg_hba.conf file of the source database. 2. Run select pg_reload_conf(); 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"> • 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. 	

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 full synchronization task contains three phases. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 3-99 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 associations are as follows: tables associated with primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by partitioned child 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. - 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.
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.
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. • 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.

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	<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.

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 3-93 Synchronization task information

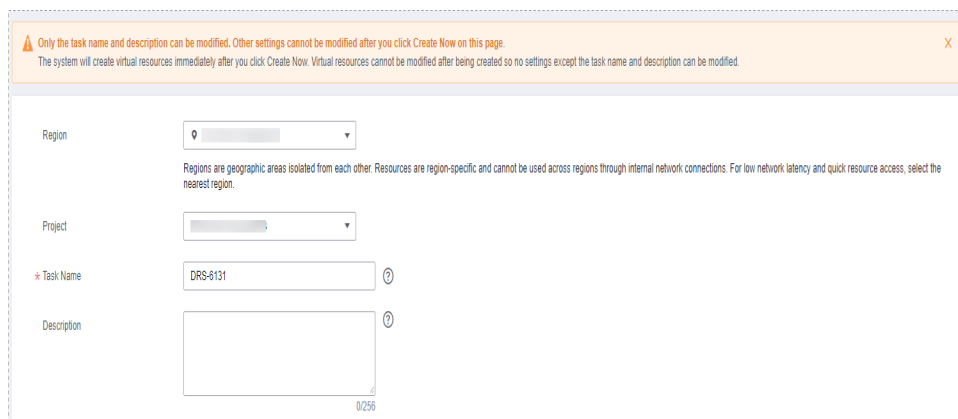


Table 3-100 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-94 Synchronization instance details

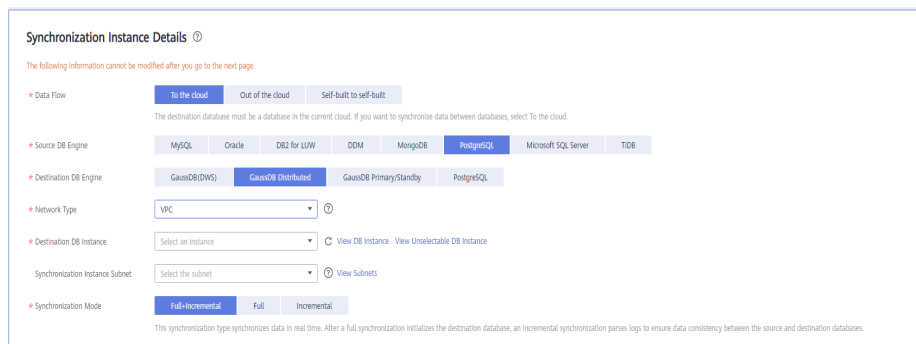


Table 3-101 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
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.
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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing.

- Enterprise Projects and Tags

Figure 3-95 Enterprise projects and tags




Table 3-102 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-96 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.

Source 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:

Table 3-103 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.

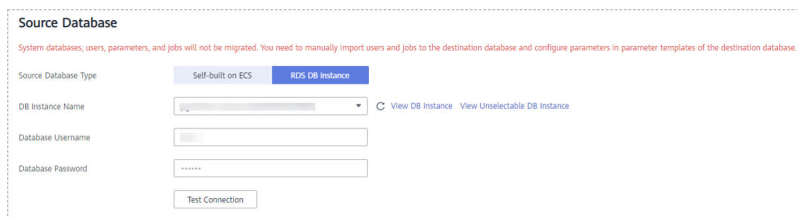
Parameter	Description
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases.

 **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 3-97 RDS DB instance - 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.

Source Database Type: Self-built on ECS RDS DB Instance

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

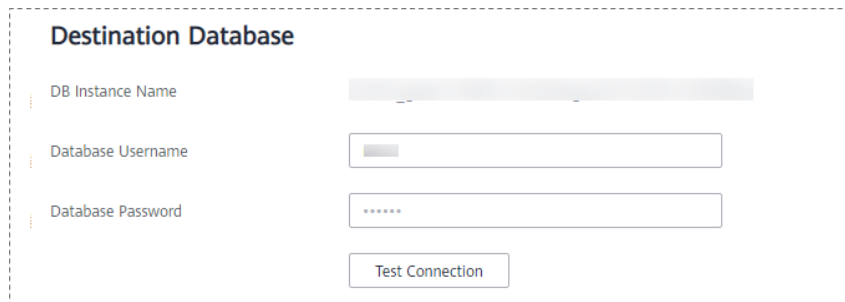
Database Username:

Database Password:

Table 3-104 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 3-98 Destination database information



Destination Database

DB Instance Name:

Database Username:

Database Password:

Table 3-105 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 3-99 Synchronization mode

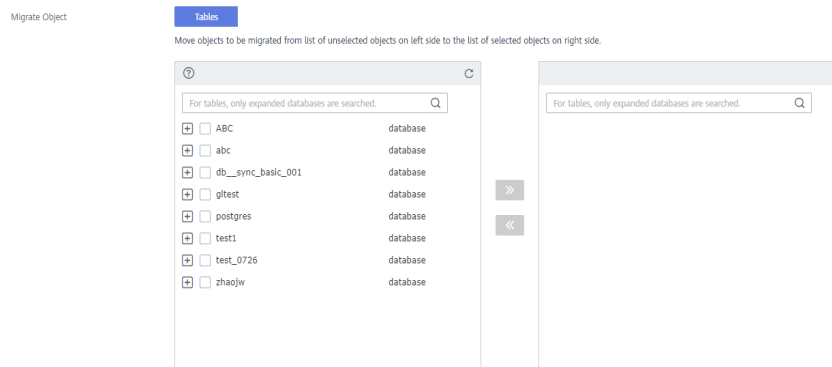



Table 3-106 Synchronization Object

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports 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 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. <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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

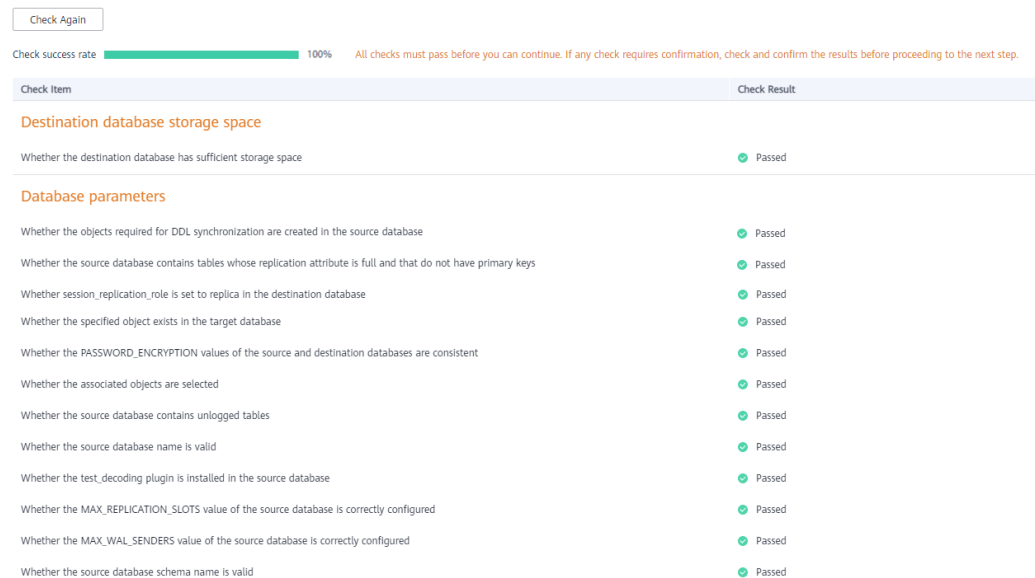
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-100 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-101 Task startup settings

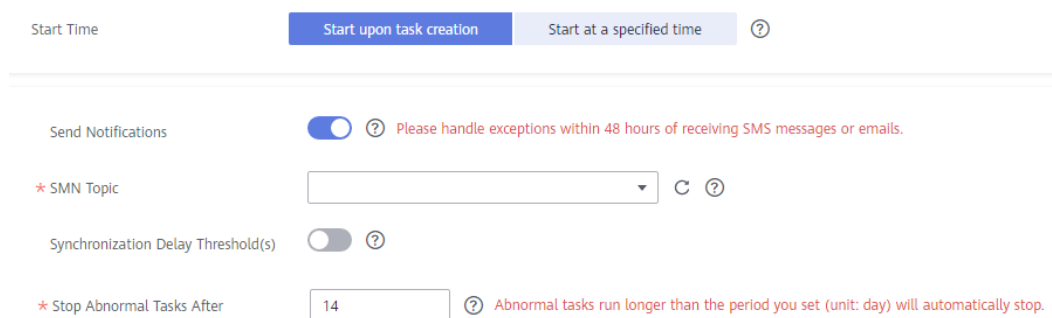



Table 3-107 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.11 From Oracle to MySQL

Supported Source and Destination Databases

Table 3-108 Supported databases

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

Supported Synchronization Objects

Table 3-109 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-109 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> • Object level: table level, object file import, and database level. • Supported synchronization objects: <ul style="list-style-type: none"> - 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. - In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. - In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. - 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. - Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by primary or foreign keys - 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 expressions of functions cannot be synchronized. - Temporary tables in the source database cannot be synchronized. - An empty source database cannot be synchronized. - If the table contains only LOB columns, data inconsistency 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.

Type	Precautions
	<p>NOTE</p> <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.

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 3-110](#). 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.

Table 3-110 Database account permission

Type	Full	Full+Incremental
Source database user	CREATE SESSION, SELECT ANY DICTIONARY, and SELECT (GRANT SELECT <userName.tbName> to drsUser) for single tables	<ul style="list-style-type: none"> • For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). • 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. • For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. • During incremental synchronization, enable PK, UK, 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

Type	Full	Full+Incremental
		<p>synchronization, ensure that the preceding settings are always enabled.</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 permission to parse logs.
Destination database user	The user must have the SELECT, INSERT, CREATE, DROP, UPDATE, ALTER, DELETE and INDEX permissions.	

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 3-111 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15. ● Source database object requirements: <ul style="list-style-type: none"> - The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibttmp1. - 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: 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

Type	Restrictions
	<p>exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle.</p> <ul style="list-style-type: none"> - 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. For detailed length requirements, see Index Length Description. - The Default User statement is not supported in MySQL. <ul style="list-style-type: none"> ● Destination database parameter requirements: <ul style="list-style-type: none"> - 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. ● Destination database object requirements: <ul style="list-style-type: none"> - 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. - 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

Type	Restrictions
	<p data-bbox="663 297 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="592 443 794 472">● Other notes: <ul style="list-style-type: none"> <li data-bbox="628 488 1366 551">– If the data types are incompatible, the synchronization may fail. <li data-bbox="628 566 1401 692">– 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. <li data-bbox="628 707 1422 875">– 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. <li data-bbox="628 891 1422 1050">– 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. <li data-bbox="628 1066 1353 1160">– After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4. <li data-bbox="628 1176 1401 1341">– 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. <li data-bbox="628 1357 1374 1482">– 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="628 1498 1426 1765">– 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="628 1780 1406 1946">– 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

Type	Restrictions
	<p>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.</p> <ul style="list-style-type: none"> - 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 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 full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. - During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated.
Full synchronization	<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 synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During synchronization, do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● 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.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none"> ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During synchronization, do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● 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 synchronization, 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 synchronization, some DDL operations are supported. <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. - Database-level synchronization supports create table. (Table definitions cannot contain functions.) - The object in 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. ● 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. ● Table names are converted to lowercase letters after the tables are 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. ● When editing the task to add a new table, ensure that transactions of the new table have been committed.

Type	Restrictions
	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.

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

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 3-102 Synchronization task information

Table 3-112 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-103 Synchronization instance details

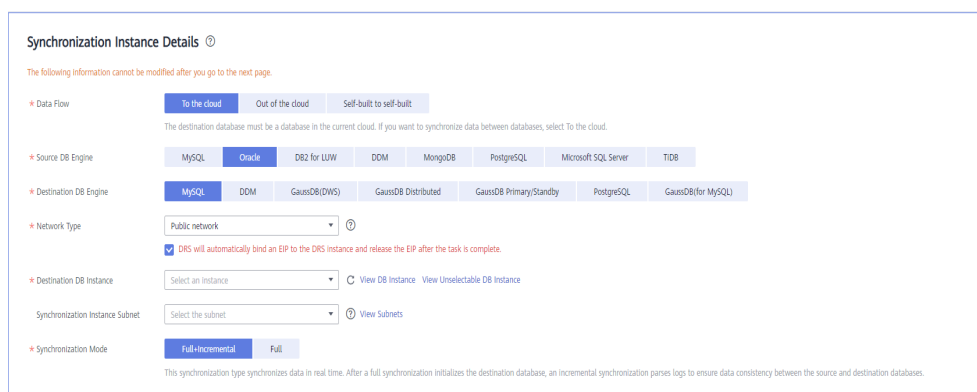


Table 3-113 Synchronization instance information

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
Destination DB Instance	The 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 Subnet 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 ensure that there are available IP addresses. 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. 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.

- DRS instance specifications

Figure 3-104 Specifications



Table 3-114 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-105 Enterprise projects and tags




Table 3-115 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-106 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

✔ Test successful

Table 3-116 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>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 not enabled, 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 3-107 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-117 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.

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 not enabled, 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 3-108 Synchronization mode

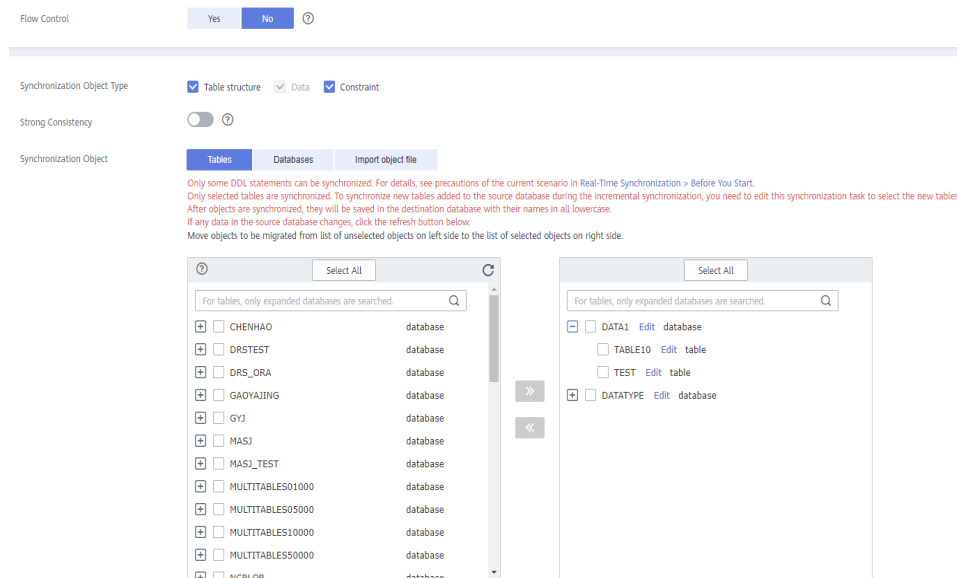
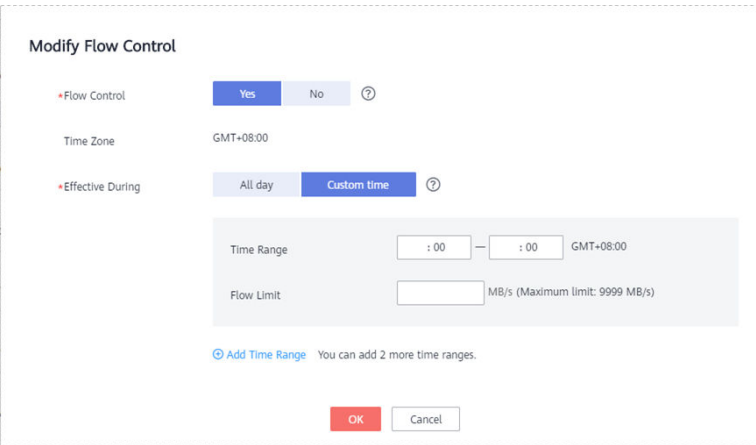



Table 3-118 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-109 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>Available options: Table structure, Data, and Constraint</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>
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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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, filter the data to be synchronized and click **Next**. For details, 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 [Checking Whether the Source Database Is Connected](#) 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-110 Task startup settings

The screenshot shows the 'Task startup settings' configuration page. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below this, there are four main settings:


- Send Notifications:** A toggle switch is turned on. A red note says 'Please handle exceptions within 48 hours of receiving SMS messages or emails.'
- SMN Topic:** A dropdown menu is empty. There are refresh and help icons.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the number '14'. A red note says 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 3-119 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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
Synchronization Delay Threshold	<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 Notification.
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> <p>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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.12 From Oracle to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-120 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	<ul style="list-style-type: none"> • GaussDB(for MySQL)

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-121 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+Incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT,

Type	Restrictions
	<p>UPDATE, ALTER, INDEX, EVENT, RELOAD, CREATE VIEW, CREATE ROUTINE, and TRIGGER.</p>
Synchronization object	<ul style="list-style-type: none"> ● 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. ● In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. ● 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. ● 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 the table contains only LOB columns, data inconsistency 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 expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized.

Type	Restrictions
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) of a single data record. ● The names of databases and tables cannot contain non-ASCII characters or special characters .><\\,?!" ● 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. ● 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. ● Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252,

Type	Restrictions
	<p>WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.</p>
Destination database	<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.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by primary or foreign keys ● The time zone settings of the source and destination database must be the same. ● If the data types are incompatible, the synchronization may 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 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. ● 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. ● 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

Type	Restrictions
	<p>contain tables with the same name but different letter cases. Otherwise, the synchronization will fail.</p> <ul style="list-style-type: none"> ● 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. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● 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, if the char field in the source database or destination database table is changed to varchar, it is padded with extra spaces. In this case, data inconsistency may occur, especially for Oracle and MySQL databases. ● 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 full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, some DDL operations are supported. <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. – Database-level synchronization supports create table. (Table definitions cannot contain functions.) – The object in DDL cannot be the keyword of the destination database, such as index or where.

Type	Restrictions
	<ul style="list-style-type: none"> - Incremental DDL operations do not support special characters such as full-width and Chinese characters. • 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.

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 3-111 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "Task Name" (a text input field containing "DRS-6131"), and "Description" (a text area). A note below the Region field states: "Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region." At the bottom right of the form, there is a character count "0/256".

Table 3-122 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-112 Synchronization instance details

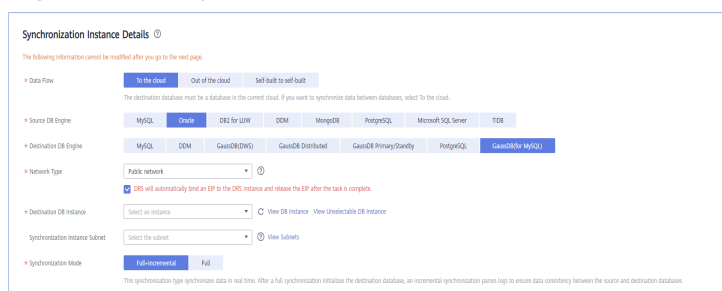


Table 3-123 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
Destination 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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

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. 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.

- DRS instance specifications

Figure 3-113 Specifications



Table 3-124 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-114 Enterprise projects and tags




Table 3-125 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-115 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

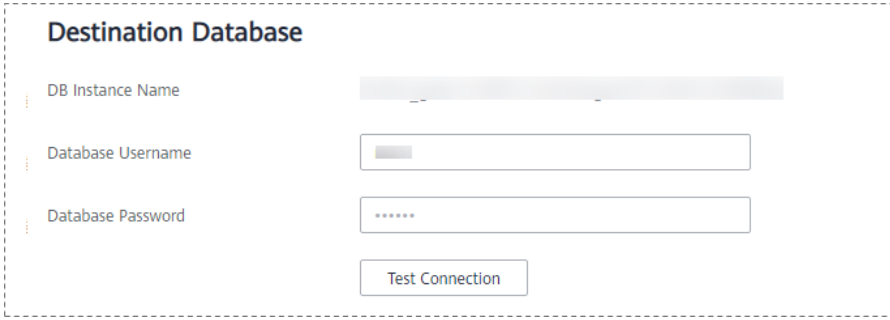
Table 3-126 Source database settings

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 not enabled, 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 3-116 Destination database information



Destination Database

DB Instance Name

Database Username

Database Password

Table 3-127 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 3-117 Synchronization mode

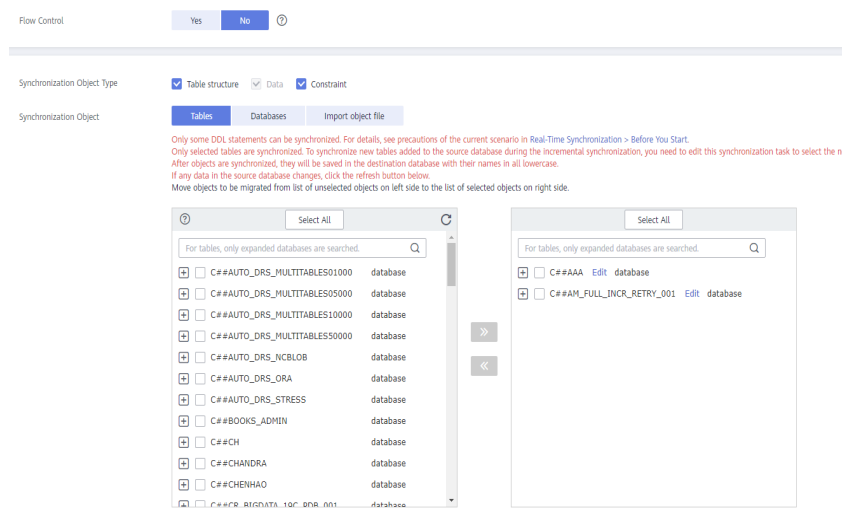
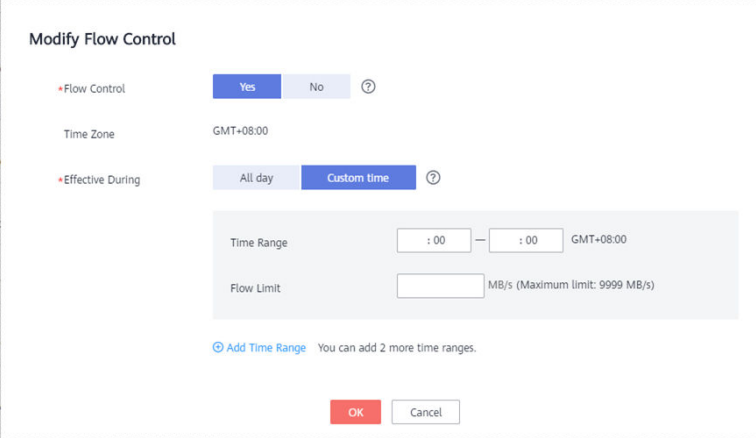



Table 3-128 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-118 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>Available options: Table structure, Data, and Constraint</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.
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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 3-119 Full synchronization

Full Synchronization Settings

★ Sharding Mode Auto Do not shard Shard ?

★ Concurrent Export Tasks - 8 + ?

★ Concurrent Import Tasks - 8 + ?

Table 3-129 Full synchronization settings

Parameter	Description	Default Value
Sharding Mode	<p>You can set this parameter to Auto, Do not shard, or Shard based on the site requirements.</p> <ul style="list-style-type: none"> • A table without a primary key is not sharded. A table with less than 520,000 rows is not sharded. Partitioned tables are sharded by partition. • If Do not shard is selected, all tables are not sharded except partitioned tables that have primary keys with more than 520,000 rows. • If Shard is selected, tables that have a primary key and more than 520,000 rows will be sharded, and other tables will not be sharded. • If Auto is selected, tables without primary keys and heap tables are not sharded. Other tables with more than 520,000 rows are sharded. 	Auto
Concurrent Export Tasks	The number of threads to be exported. The value ranges from 1 to 16.	8
Concurrent Import Tasks	The number of threads to be imported. The value ranges from 1 to 16.	8

Figure 3-120 Incremental synchronization

Incremental Data Capture Settings

★ Concurrent Log Capture Tasks ?

Table 3-130 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 [Checking Whether the Source Database Is Connected](#) 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-121 Task startup settings


The screenshot shows the 'Task startup settings' configuration page. At the top, there are two radio buttons for 'Start Time': 'Start upon task creation' (selected) and 'Start at a specified time'. Below this, there are several settings: 'Send Notifications' is a checked toggle switch with a help icon and a note: 'Please handle exceptions within 48 hours of receiving SMS messages or emails.' 'SMN Topic' is a dropdown menu with a refresh icon and a help icon. 'Synchronization Delay Threshold(s)' is an unchecked toggle switch with a help icon. 'Stop Abnormal Tasks After' is a text input field containing '14' with a help icon and a note: 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 3-131 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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
Synchronization Delay Threshold	<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 Notification.
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> <p>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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.13 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-132 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	GaussDB primary/standby

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-133 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+incremental synchronization and incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● The destination database must have the following permissions:

Type	Restrictions
	<ul style="list-style-type: none"><li data-bbox="627 300 1390 427">- 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.<li data-bbox="627 443 1409 571">- 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.<li data-bbox="627 586 1426 745">- 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, 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, xml, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, xml, sdo_geometry, urowid, interval, and user-defined types. ● The maximum precision supported by timestamp and interval day to second is 6. ● 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. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, and function indexes. ● 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. ● 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 data 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 the table contains only LOB columns, data inconsistency 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.

Type	Restrictions
	<ul style="list-style-type: none"> • 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 expressions of functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • 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 due to time zone conversion.
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 following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.

Type	Restrictions
Destination database	<ul style="list-style-type: none">● The destination database must be a GaussDB primary/standby instance on the current cloud.● 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:<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.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• 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.• 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 and value comparison will be unavailable. For details about the data types supported by the primary key, see Mapping Data Types.• 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.• 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.• If the PDB database is used for synchronization, all PDBs must be enabled during 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.• 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.• 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

Type	Restrictions
	<p>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.</p> <ul style="list-style-type: none">• 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.• During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• 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.• 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 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.• 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.• 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.

Type	Restrictions
	<ul style="list-style-type: none"> ● During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● When a table structure is fully synchronized, only default value constraints of the character string or number type are supported. Default value constraints of the function and sequence types are not supported. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, some DDL operations are supported. <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. – Database-level synchronization supports create table. (Table definitions cannot contain functions.) ● 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. ● Full and incremental synchronization does not support the STRING extended data type (EXTENDED DATA TYPE) of Oracle 12c or later. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the source database uses XStream for log reading, you need to insert data during incremental migration. Otherwise, the incremental migration delay keeps increasing. ● When DRS reads incremental Oracle logs in XStream mode, an Oracle XStream outbound is created on the source database. If another XStream outbound that reports an error already exists on the source database, the new XStream outbound created by DRS may be affected. As a result, data cannot be read. You can run the following command in the source database to check whether there are other XStream outbound errors: <pre style="background-color: #f0f0f0; padding: 5px;">select o.server_name, c.capture_name, c.error_message from SYS.ALL_XSTREAM_OUTBOUND O, SYS.ALL_CAPTURE C where o.capture_name = c.capture_name</pre>

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.

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 3-122 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-6131 [Text input]

Description: [Text area]

0/256

Table 3-134 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-123 Synchronization instance details




Table 3-135 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	The public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
Destination DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full, Incremental , and Full +Incremental .

- Enterprise Projects and Tags

Figure 3-124 Enterprise projects and tags




Table 3-136 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-125 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 3-137 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 not enabled, your data may be at risk.

Figure 3-126 Destination database information

The screenshot shows a form titled "Destination Database" with the following elements:

- DB Instance Name:** A text input field with a grey background.
- Database Username:** A text input field with a grey background.
- Database Password:** A text input field with a grey background and masked characters (*****).
- Test Connection:** A button located below the password field.

Table 3-138 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. 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 3-127 Synchronization mode

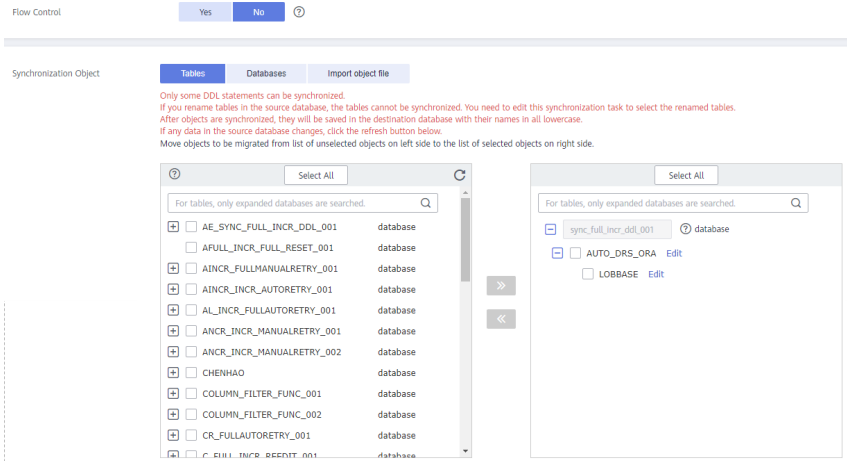
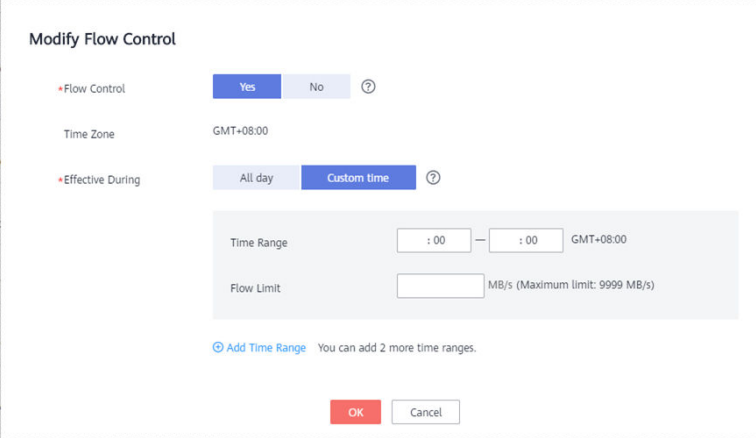



Table 3-139 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-128 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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, view the parameters you set for full synchronization or incremental synchronization in 2. The default parameter values cannot be modified. Click **Next**.


Figure 3-129 Full synchronization

Full Synchronization Settings


* Synchronization Object Type Table structure Data Constraint

* Stream Mode 

* Concurrent Export Tasks 

* Concurrent Import Tasks 

* Import Mode COPY INSERT
 The copy mode improves write efficiency.

* Rows per Shard 


Synchronization Point 

Table 3-140 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Available options: Table structure , Data , and Constraint Data is mandatory.	All options
Stream Mode	In this mode, data is written to the destination database by shard. Data is submitted once for each shard. Otherwise, The data is submitted once every 16 MB of data is written to the destination database.	Enabled
Concurrent Export Tasks	The number of threads that can be started to export data from the Oracle database.	8
Concurrent Import Tasks	The number of threads that can be started to concurrently write data to GaussDB.	8
Import Mode	<ul style="list-style-type: none">• COPY With the COPY interface, writing efficiency is relatively high, but task errors may occur due to data conflicts.• INSERT Writing data using the INSERT statement can avoid data conflicts but the writing efficiency is low.	COPY
Rows per Shard	Export tables in shards to improve the efficiency of real-time synchronization. <ul style="list-style-type: none">• 0 All tables are not sharded, and each table is synchronized as a whole.• Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded.• If the table is a partitioned table, each partition is synchronized as a shard instead of being sharded based on the value you set.• If a table has no primary key or unique index, the table is not sharded. In this case, this parameter is meaningless.• If a table is not analyzed in a timely manner, the value of NUM_ROWS in the ALL_TABLES view is empty, and the table is not sharded.	520000

Parameter	Description	Default Value
Synchronization Point	Set this parameter to the current system change number (SCN). If you set this parameter to a future SCN or the SCN range where deleted archive logs are located, the task will fail.	The default value is the SCN generated when the task is started.

Figure 3-130 Incremental synchronization

Incremental Data Capture Settings

* Concurrent Log Capture Tasks ?

Capture Start Point

If this parameter is not specified, the start point of the current database is used. (In Oracle database, SCN is used.)

Incremental Replay Settings

* Concurrent Replay Tasks ?

* Replay Policy Automated Manual
After the task is started, the replay component is started automatically.

* Conflict Policy Overwrite Report error Ignore
The destination database data is overwritten by the new data.

Replay Start Point

If this parameter is not specified, the replay start point is the same as the capture start point by default.

Table 3-141 Incremental synchronization settings

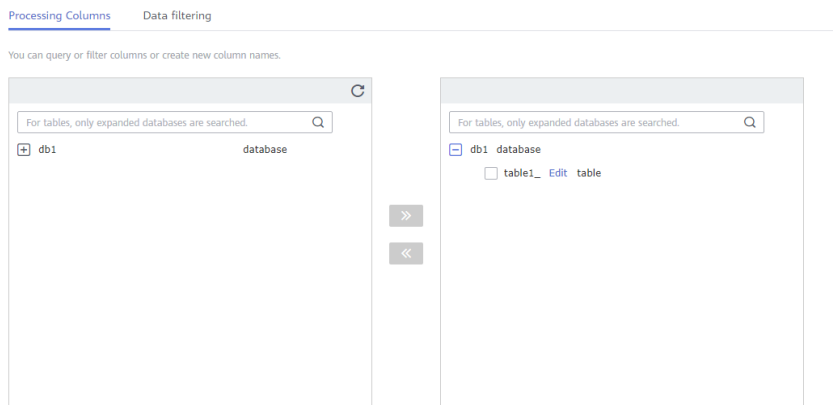
Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads for reading Oracle logs. The value ranges from 1 to 16.	1
Capture Start Point	Specifies the SCN for starting the capture. SCNs are designed to meet service requirements. It consists of a start SCN for capturing and a valid SCN. For details, see the SCN concepts of Oracle.	If this parameter is left blank, the current SCN of the database is used as the start point by default.

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
Replay Policy	<ul style="list-style-type: none"> • Automated After the task is started, the replay component is started automatically. • Manual After a task is started, the replay component needs to be started manually. 	Auto
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
Replay Start Point	Specifies the SCN where the replay starts.	If this parameter is not specified, the replay start point is the same as the capture start point by default.

Step 6 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 3-131 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-132 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-133 Task startup settings

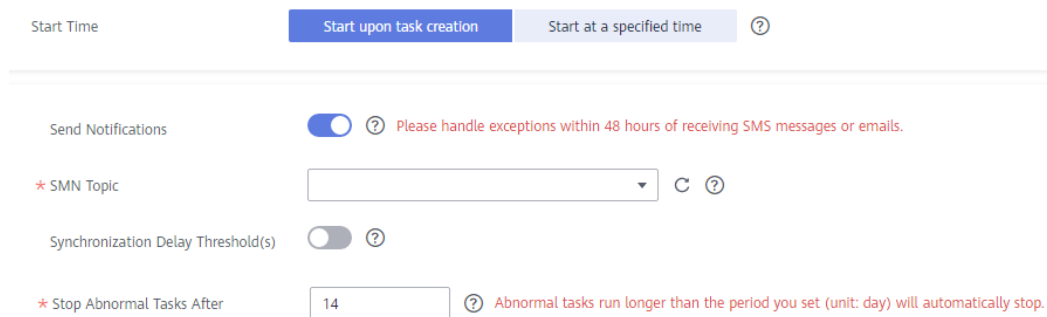



Table 3-142 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.14 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 3-143 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	GaussDB distributed

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-144 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+incremental synchronization and incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● The destination database must have the following permissions:

Type	Restrictions
	<ul style="list-style-type: none"><li data-bbox="627 300 1390 427">- 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.<li data-bbox="627 443 1409 571">- 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.<li data-bbox="627 586 1426 745">- 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, 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, xmltype, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: XMLTYPE, bfile, xmltype, sdo_geometry, urowid, interval year to month, interval day to second, and user-defined types. ● 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 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. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, and function indexes. ● The maximum precision supported by timestamp and interval day to second is 6. ● 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. ● For the TIMESTAMP WITH TIME ZONE data type, the data 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 the table contains only LOB columns, data inconsistency 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

Type	Restrictions
	<p>the processing of the empty string in the destination database.</p> <ul style="list-style-type: none"> ● 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 expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● 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 due to time zone conversion.
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 following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.

Type	Restrictions
Destination database	<ul style="list-style-type: none">● The destination database must be a GaussDB distributed instance on the current cloud.● 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:<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.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• 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.• 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 and value comparison will be unavailable. For details about the data types supported by the primary key, see Mapping Data Types.• 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.• 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.• If the PDB database is used for synchronization, all PDBs must be enabled during 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.• 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.• 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

Type	Restrictions
	<p>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.</p> <ul style="list-style-type: none"> ● 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. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● 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. ● 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 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. ● 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, if the destination database has a constraint with the same name as the source database, the table fails to be created.

Type	Restrictions
	<ul style="list-style-type: none"> ● During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● When a table structure is fully synchronized, only default value constraints of the character string or number type are supported. Default value constraints of the function and sequence types are not supported. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, some DDL operations are supported. <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. – Database-level synchronization supports create table. (Table definitions cannot contain functions.) ● 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. ● Full and incremental synchronization does not support the STRING extended data type (EXTENDED DATA TYPE) of Oracle 12c or later. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the source database uses XStream for log reading, you need to insert data during incremental migration. Otherwise, the incremental migration delay keeps increasing. ● When DRS reads incremental Oracle logs in XStream mode, an Oracle XStream outbound is created on the source database. If another XStream outbound that reports an error already exists on the source database, the new XStream outbound created by DRS may be affected. As a result, data cannot be read. You can run the following command in the source database to check whether there are other XStream outbound errors: <pre style="background-color: #f0f0f0; padding: 5px;">select o.server_name, c.capture_name, c.error_message from SYS.ALL_XSTREAM_OUTBOUND O, SYS.ALL_CAPTURE C where o.capture_name = c.capture_name</pre>

Type	Restrictions
	<ul style="list-style-type: none"> • If the destination database is a GaussDB distributed 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 the distribution column. • 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.

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 3-134 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is an orange warning banner with a triangle icon and the text: "Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page." Below the banner, there is a sub-note: "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." The form contains the following fields:

- Region:** A dropdown menu with a small 'Q' icon on the left.
- Project:** A dropdown menu.
- * Task Name:** A text input field containing "DRS-6131" and a help icon (circle with 'i').
- Description:** A text area with a help icon (circle with 'i').

 At the bottom right of the form, there is a small "0/256" character count.

Table 3-145 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-135 Synchronization instance details

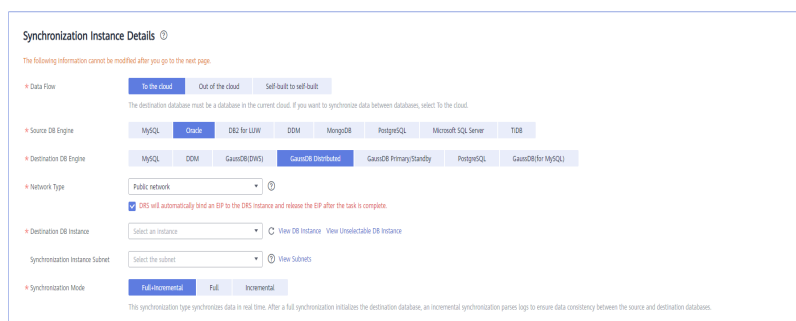


Table 3-146 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
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
Destination DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	Available options: Full , Incremental , and Full +Incremental .

- Enterprise Projects and Tags

Figure 3-136 Enterprise projects and tags

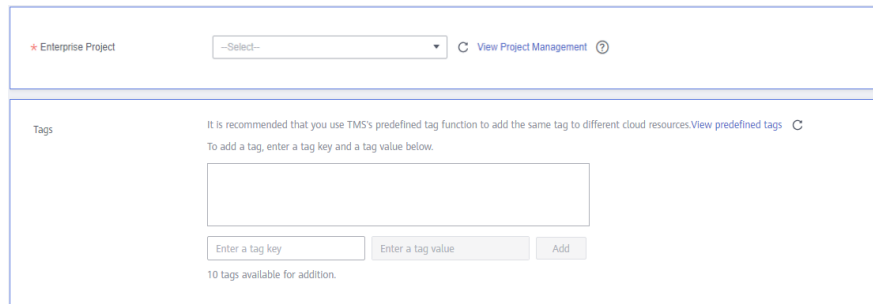


Table 3-147 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

- 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 3-137 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 3-148 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 not enabled, your data may be at risk.

Figure 3-138 Destination database information

The screenshot shows a form titled "Destination Database" with the following elements:

- DB Instance Name:** A text input field with a grey background.
- Database Username:** A text input field with a grey background.
- Database Password:** A text input field with a grey background and masked characters (*****).
- Test Connection:** A button located below the password field.

Table 3-149 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. 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 3-139 Synchronization mode

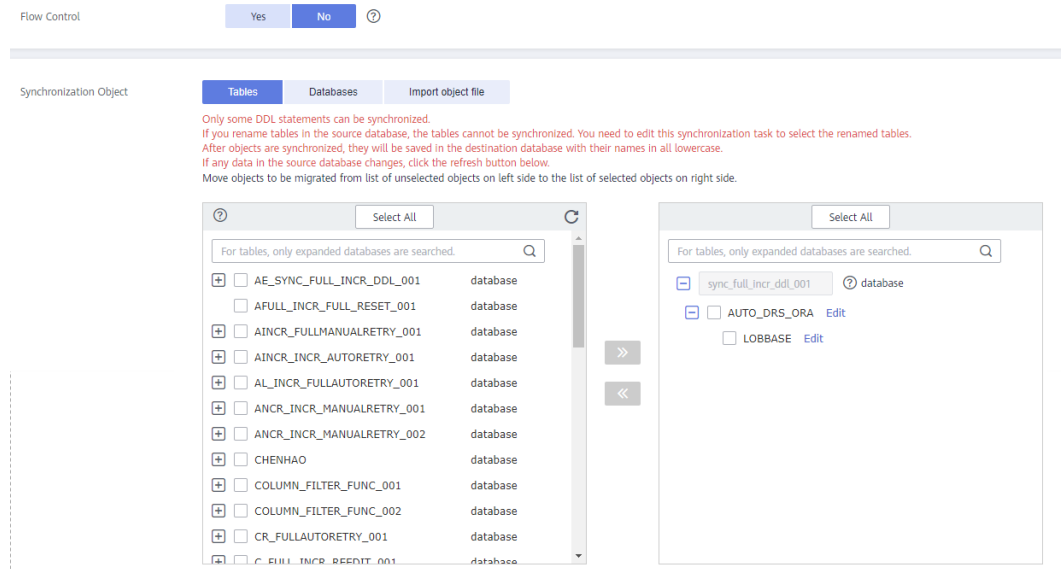
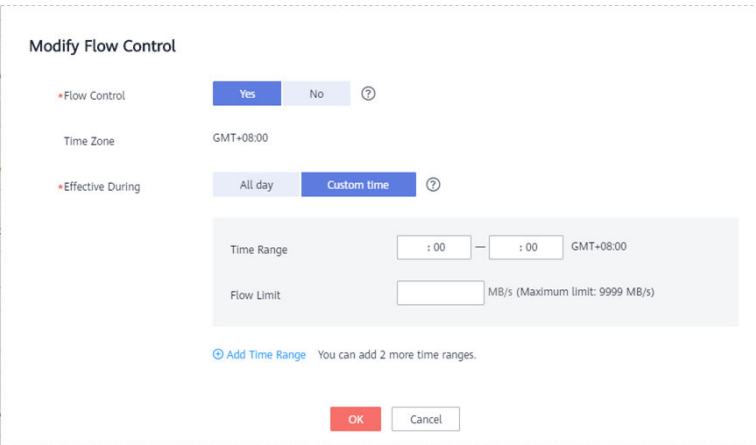



Table 3-150 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-140 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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, view the parameters you set for full synchronization or incremental synchronization in [2](#). The default parameter values cannot be modified. Click **Next**.


Figure 3-141 Full synchronization

Full Synchronization Settings


* Synchronization Object Type Table structure Data Constraint

* Stream Mode 

* Concurrent Export Tasks 

* Concurrent Import Tasks 

* Import Mode COPY INSERT
 The copy mode improves write efficiency.

* Rows per Shard 


Synchronization Point 

Table 3-151 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Available options: Table structure , Data , and Constraint Data is mandatory.	All options
Stream Mode	In this mode, data is written to the destination database by shard. Data is submitted once for each shard. Otherwise, The data is submitted once every 16 MB of data is written to the destination database.	Enabled
Concurrent Export Tasks	The number of threads that can be started to export data from the Oracle database.	8
Concurrent Import Tasks	The number of threads that can be started to concurrently write data to GaussDB.	8
Import Mode	<ul style="list-style-type: none"> • COPY With the COPY interface, writing efficiency is relatively high, but task errors may occur due to data conflicts. • INSERT Writing data using the INSERT statement can avoid data conflicts but the writing efficiency is low. 	COPY
Rows per Shard	<p>Export tables in shards to improve the efficiency of real-time synchronization.</p> <ul style="list-style-type: none"> • 0 All tables are not sharded, and each table is synchronized as a whole. • Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. • If the table is a partitioned table, each partition is synchronized as a shard instead of being sharded based on the value you set. • If a table has no primary key or unique index, the table is not sharded. In this case, this parameter is meaningless. • If a table is not analyzed in a timely manner, the value of NUM_ROWS in the ALL_TABLES view is empty, and the table is not sharded. 	520000

Parameter	Description	Default Value
Synchronization Point	Set this parameter to the current system change number (SCN). If you set this parameter to a future SCN or the SCN range where deleted archive logs are located, the task will fail.	The default value is the SCN generated when the task is started.

Figure 3-142 Incremental synchronization

Incremental Data Capture Settings

* Concurrent Log Capture Tasks ?

Capture Start Point

If this parameter is not specified, the start point of the current database is used. (In Oracle database, SCN is used.)

Incremental Replay Settings

* Concurrent Replay Tasks ?

* Replay Policy Automated Manual

After the task is started, the replay component is started automatically.

* Conflict Policy Overwrite Report error Ignore

The destination database data is overwritten by the new data.

Replay Start Point

If this parameter is not specified, the replay start point is the same as the capture start point by default.

Table 3-152 Incremental synchronization settings

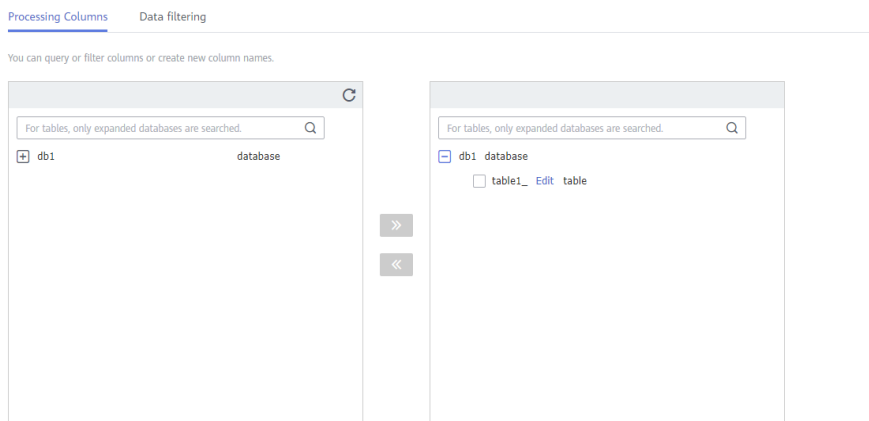
Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads for reading Oracle logs. The value ranges from 1 to 16.	1
Capture Start Point	Specifies the SCN for starting the capture. SCNs are designed to meet service requirements. It consists of a start SCN for capturing and a valid SCN. For details, see the SCN concepts of Oracle.	If this parameter is left blank, the current SCN of the database is used as the start point by default.

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
Replay Policy	<ul style="list-style-type: none"> • Automated After the task is started, the replay component is started automatically. • Manual After a task is started, the replay component needs to be started manually. 	Auto
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
Replay Start Point	Specifies the SCN where the replay starts.	If this parameter is not specified, the replay start point is the same as the capture start point by default.

Step 6 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 3-143 Processing data



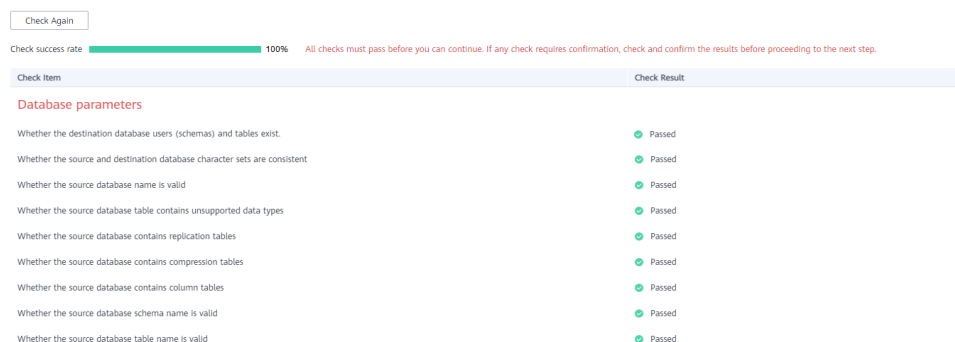
Step 7 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-144 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-145 Task startup settings





Table 3-153 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.15 From Oracle to DDM

Supported Source and Destination Databases

Table 3-154 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	<ul style="list-style-type: none"> • DDM instances

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-155 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+Incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD and CREATE VIEW.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● DDL operations cannot be synchronized during incremental synchronization. ● In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. ● 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. ● The source database data can be synchronized, but the source table structure and other objects cannot be synchronized. ● Create table structures and indexes in the destination database that corresponds to the schema of the source database. 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 the table contains only LOB columns, data inconsistency 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 expressions of functions cannot be synchronized.

Type	Restrictions
	<ul style="list-style-type: none"> Temporary tables in the source database cannot be synchronized.
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 (<code>_</code>) 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. Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.

Type	Restrictions
Destination database	<ul style="list-style-type: none"> ● 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 end. ● 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.

Type	Restrictions
Operations	<ul style="list-style-type: none"> ● Table names are converted to lowercase letters after the tables are 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. ● 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) 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 data types are incompatible, the synchronization may fail.

Type	Restrictions
	<ul style="list-style-type: none"> ● 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. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During synchronization, DDL operations cannot be performed on the source database. ● 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 synchronization, writing data to the destination databases is not allowed. Otherwise, data inconsistency may occur. ● 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 full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● 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.

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.

3.16 From Oracle to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-156 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises databases ECS databases 	<ul style="list-style-type: none"> GaussDB(DWS) cluster

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-157 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+Incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● Destination database:

Type	Restrictions
	<ul style="list-style-type: none"><li data-bbox="627 300 1422 465">- 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.<li data-bbox="627 477 1422 539">- The INSERT, SELECT, UPDATE, and DELETE permissions are required for each table.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, indexes, and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraint, and unique index, cannot be synchronized. ● In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. ● 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 the TIMESTAMP WITH TIME ZONE data type, the data 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 the table contains only LOB columns, data inconsistency 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 expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● When you manually create a table structure, the time type in the destination database must be the same as that in the

Type	Restrictions
	<p>source database. Otherwise, time data may be inconsistent due to time zone conversion.</p>
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. ● 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 DB instance is running properly. ● 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 may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. ● 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. ● 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 when indexes are synchronized. Primary key constraints are synchronized with the table structure. ● 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 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) 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.

Type	Restrictions
	<ul style="list-style-type: none"> ● 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 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. ● Multiple source tables can be mapped to one GaussDB(DWS) table. For details, see From Oracle to GaussDB(DWS). ● During synchronization, do not delete the username, password, and permissions of the source and destination databases or change the port of the destination database. ● 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. ● 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 full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, the following DDL operations are supported:

Type	Restrictions
	<ul style="list-style-type: none"> - Table-level synchronization supports the following DDL operations: alter table add column, alter table drop column, alter table modify column, and truncate table • 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.

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 3-146 Synchronization task information

The screenshot shows a web form for creating a synchronization instance. At the top, there is an orange warning banner with a close button (X) that reads: "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." Below the banner, the form contains the following fields:

- Region:** A dropdown menu with a search icon.
- Project:** A dropdown menu.
- * Task Name:** A text input field containing "DRS-0131" and a help icon.
- Description:** A text area with a help icon.

 A small "0/255" character count is visible at the bottom right of the description field.

Table 3-158 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-147 Synchronization instance details

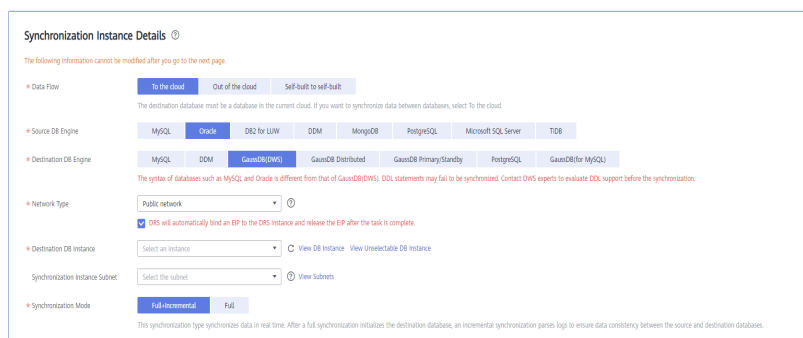


Table 3-159 Synchronization instance settings

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
Destination DB Instance	An available GaussDB(DWS) instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

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. 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.

- DRS instance specifications

Figure 3-148 Specifications



Table 3-160 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-149 Enterprise projects and tags




Table 3-161 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-150 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 3-162 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 not enabled, 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 3-151 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-163 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 3-152 Synchronization mode

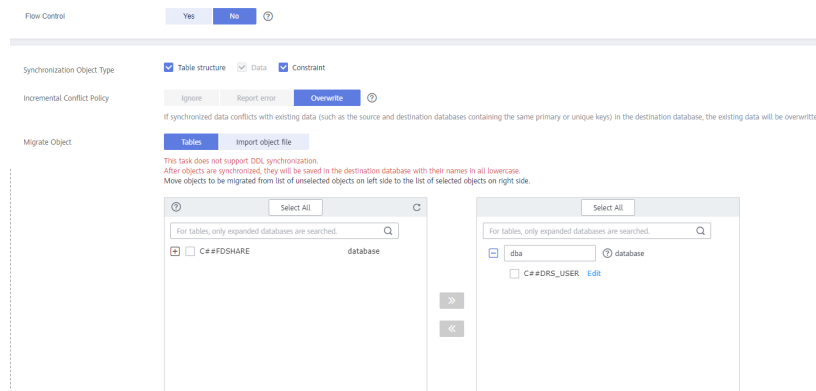
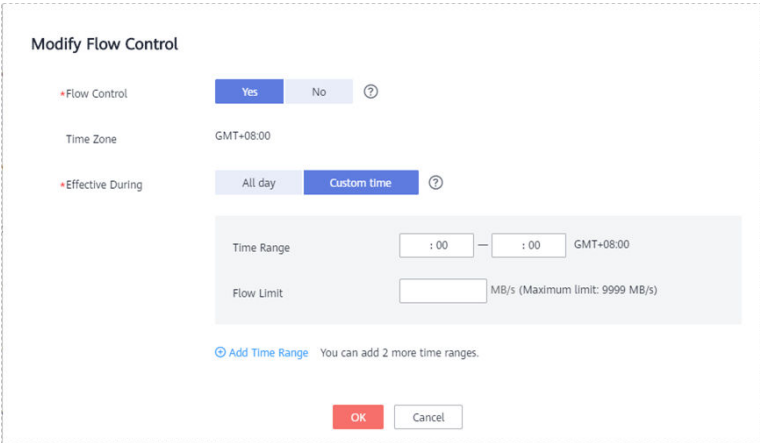



Table 3-164 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-153 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>Available options: Table structure, Data, and Constraint</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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 3-154 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
CHENHAO.AB	CHENHAO.AB	--	--	--	Add
CHENHAO.ABC	CHENHAO.ABC	--	--	--	Add
CHENHAO.DIF	CHENHAO.DIF	--	--	--	Add
CHENHAO.IUFO_MEASPU8_XSFA	CHENHAO.IUFO_MEASPU8_XSFA	--	--	--	Add
CHENHAO.IUFO_MEASPU8_XSFA	CHENHAO.IUFO_MEASPU8_XSFA	--	--	--	Add
CHENHAO.NA	CHENHAO.NA	--	--	--	Add
CHENHAO.NB	CHENHAO.NB	--	--	--	Add
CHENHAO.ONE	CHENHAO.ONE	--	--	--	Add
CHENHAO.ORE1	CHENHAO.ORE1	--	--	--	Add
CHENHAO.ORACLE1	CHENHAO.ORACLE1	--	--	--	Add

10 Total Records: 3,913 2 3 4 5 ... 392 >

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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-155 Pre-check

Check Again

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 values in the source Oracle database are out of the ranges of the MySQL database	⚠ Confirm Confirm Details
Whether the selected objects can be migrated	✔ Passed
Whether the source database contains unsupported table field types	✔ Passed
Whether tables to be migrated contain primary keys	✔ Passed
Whether the source database character set is supported	✔ Passed
Whether the source database container type is correct	✔ Passed
Whether archive logs are enabled on the source database	✔ Passed
Whether the source database name is valid	✔ Passed
Whether the supplementary log is enabled for the source database.	✔ Passed
Whether the destination database contains the same object names as those of the synchronization objects	✔ 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-156 Task startup settings

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

Send Notifications ? Please handle exceptions within 48 hours of receiving SMS messages or emails.

* SMN Topic ?

Synchronization Delay Threshold(s) ?


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

Table 3-165 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.17 From Oracle to PostgreSQL

Supported Source and Destination Databases

Table 3-166 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	<ul style="list-style-type: none"> • RDS for PostgreSQL

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.

- Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 3-167 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+Incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, indexes, and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, and virtual columns cannot be synchronized. ● In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. ● 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 the TIMESTAMP WITH TIME ZONE data type, the data 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 the table contains only LOB columns, data inconsistency 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 expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● When you manually create a table structure, the time type in the destination database must be the same as that in the

Type	Restrictions
	<p>source database. Otherwise, time data may be inconsistent due to time zone conversion.</p>
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. ● 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 DB instance is running properly. ● 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 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.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. ● 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 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) 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

Type	Restrictions
	<p>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.</p> <ul style="list-style-type: none"> ● 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. ● During synchronization, do not delete the username, password, and permissions of the source and destination databases or change the port of the destination database. ● 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. ● 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 full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. ● Only normal indexes are synchronized when indexes are synchronized. Primary key constraints are synchronized with the table structure. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, some DDL operations are supported. <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"> • The names of mapped table-level objects are case-insensitive. For example, the ABC table 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. • 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. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p>

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](#).

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 3-157 Synchronization task information

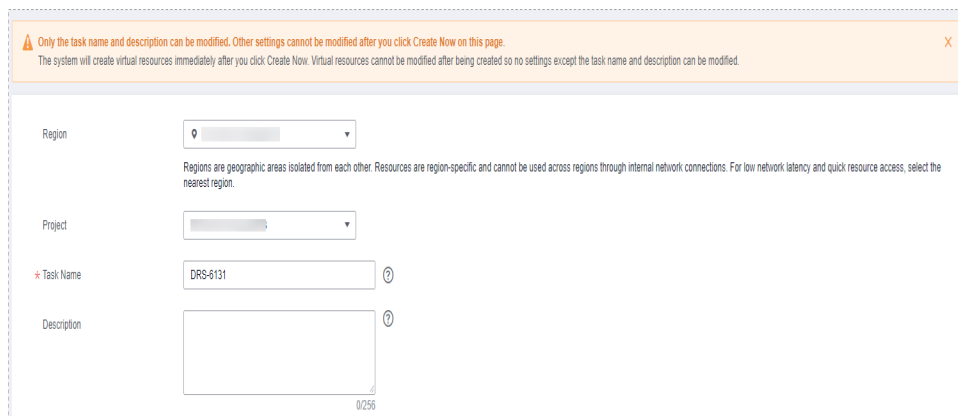


Table 3-168 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-158 Synchronization instance details




Table 3-169 Synchronization instance settings

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

Parameter	Description
Destination DB Engine	Select PostgreSQL .
Network Type	Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
Destination DB Instance	The RDS for PostgreSQL instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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. 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.

- DRS instance specifications

Figure 3-159 Specifications



Table 3-170 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-160 Enterprise projects and tags




Table 3-171 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 3-161 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

Test successful

Table 3-172 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 not enabled, 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 3-162 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 3-173 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 3-163 Synchronization mode

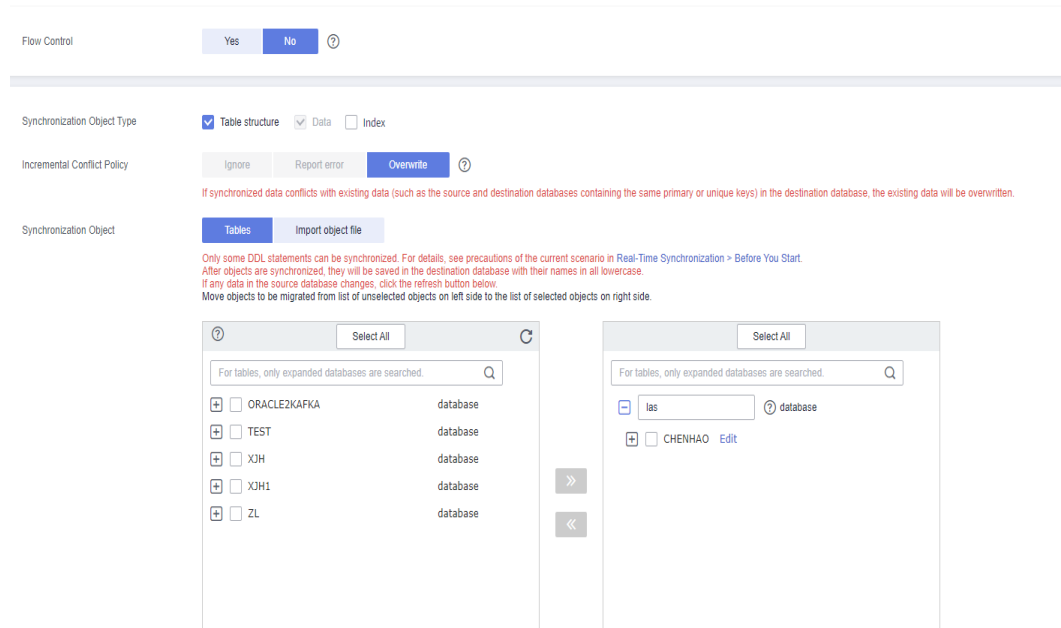
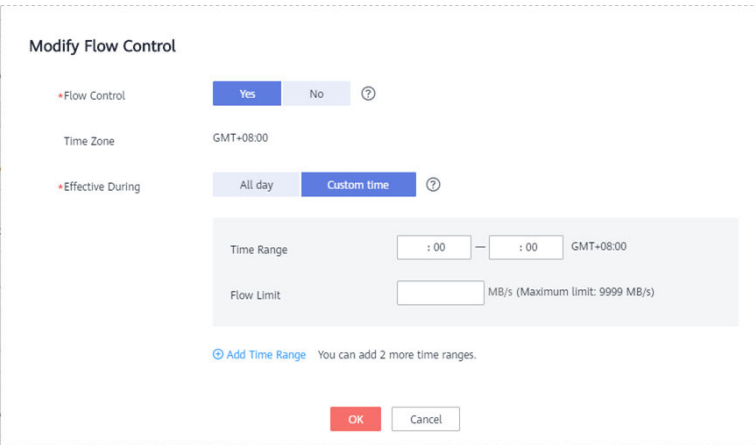



Table 3-174 Synchronization mode and object

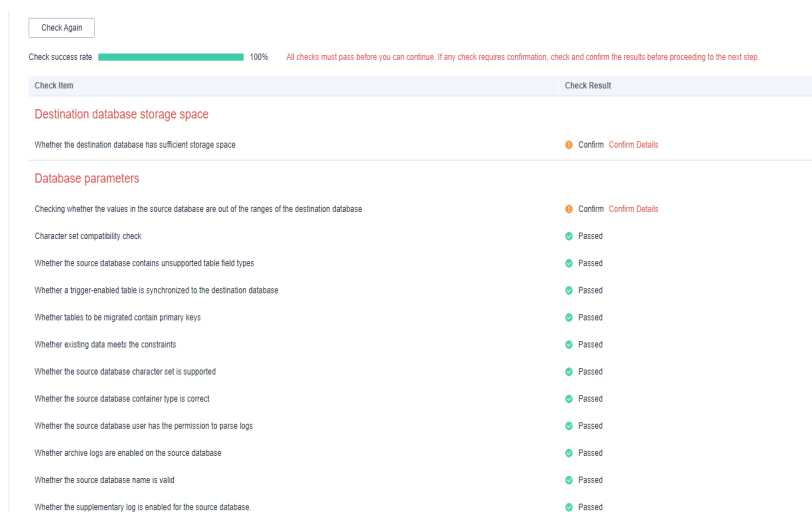
Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of 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 All day. A maximum of three 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 3-164 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>Available options: Table structure, Data, and Constraint</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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 3-165 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-166 Task startup settings

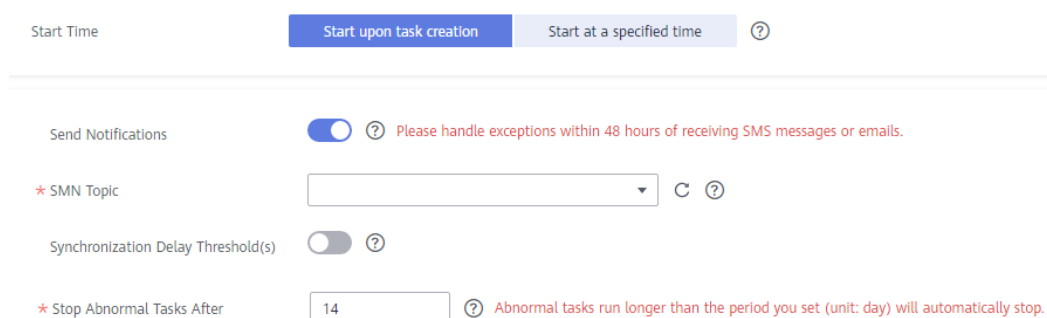



Table 3-175 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.18 From DDM to MySQL

Supported Source and Destination Databases

Table 3-176 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">DDM instances	<ul style="list-style-type: none">RDS for MySQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 3-177 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 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: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, and UPDATE. The root account of the RDS for MySQL DB instance has the preceding permissions by default.
Synchronization object	<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 timestamp. • 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.

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 source database binlog 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 database sharding middleware cannot contain the following characters: '<>\' and non-ASCII characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● Enable GTID of the source database.
Destination database	<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 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. ● 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.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a 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 destination DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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. ● 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. ● During incremental synchronization, do not perform the restoration operation on the source database. ● During incremental synchronization, 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.

Type	Restrictions
	<ul style="list-style-type: none"> - 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 '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. - 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. • During incremental synchronization, perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute.

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 3-167 Synchronization task information

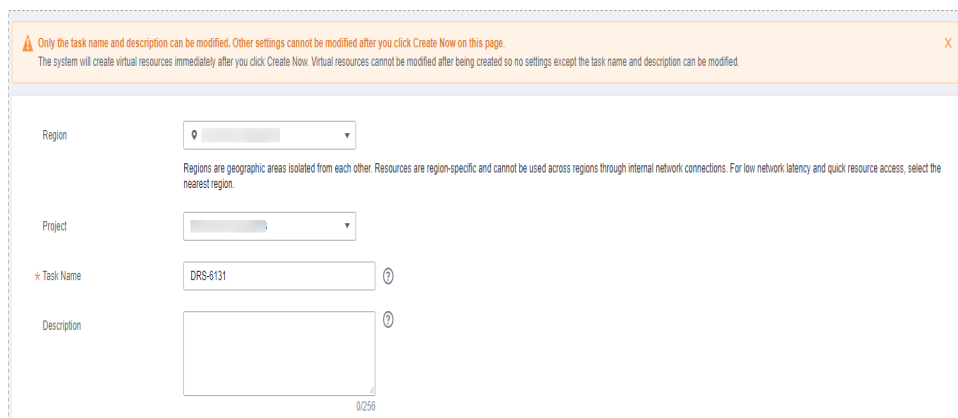


Table 3-178 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-168 Synchronization instance details

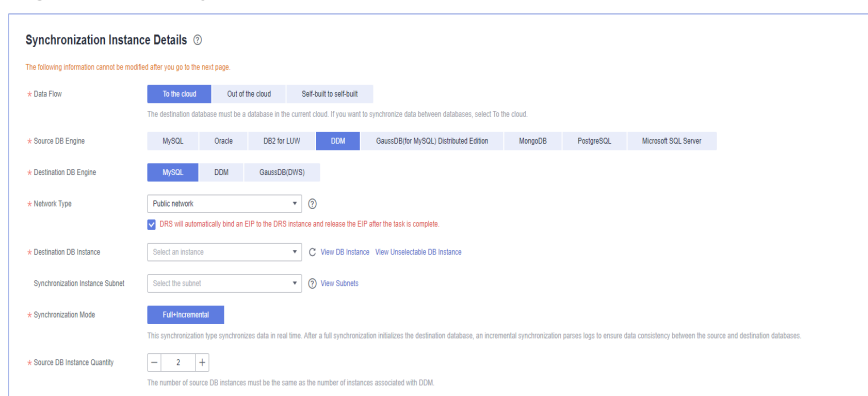


Table 3-179 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .

Parameter	Description
Source DB Engine	Select DDM .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
Destination DB Instance	The RDS DB instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	<p>– 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> <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.

- DRS instance specifications

Figure 3-169 Specifications



Table 3-180 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-170 Enterprise projects and tags

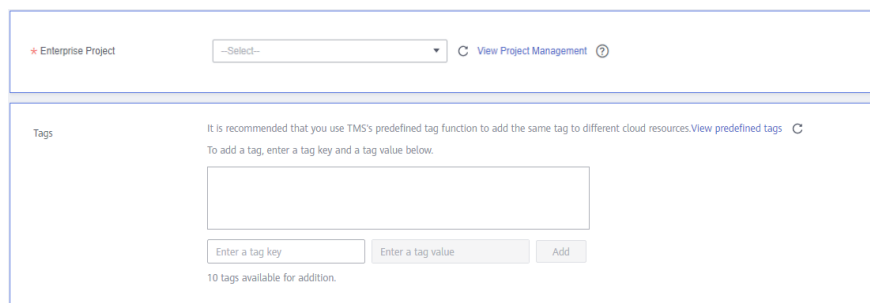


Table 3-181 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 3-171 Source database information

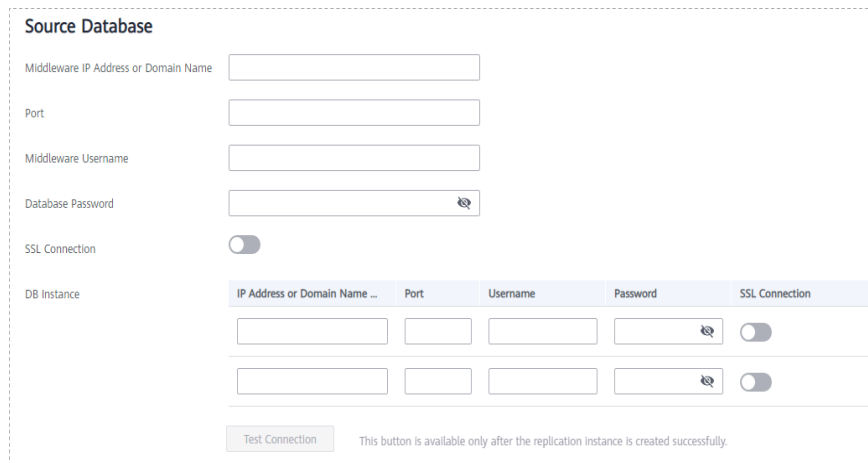


Table 3-182 Source database settings

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	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 not enabled, 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 3-172 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-183 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.
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 not enabled, 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 3-173 Synchronization mode

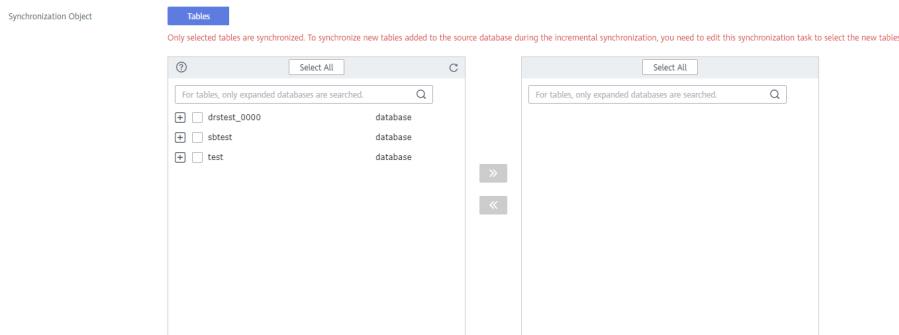



Table 3-184 Synchronization mode and object

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. ● 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

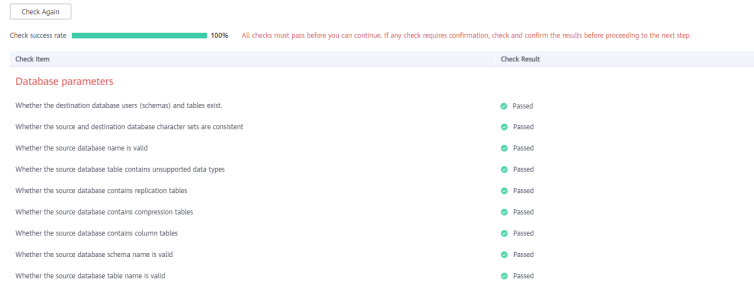
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-174 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-175 Task startup settings

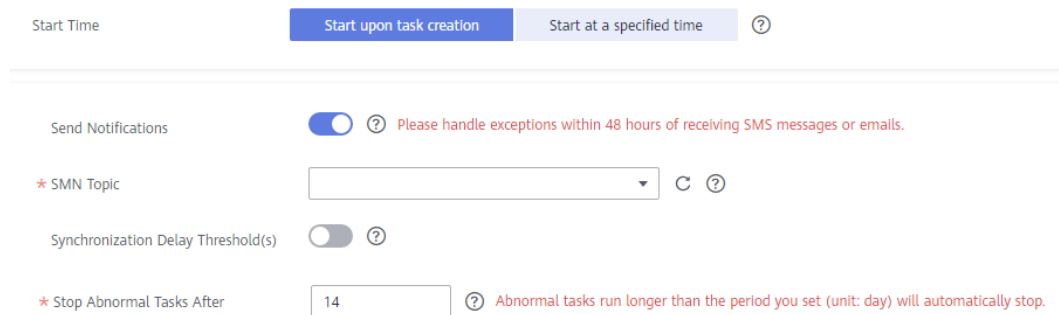



Table 3-185 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.19 From DDM to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-186 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">DDM instances	<ul style="list-style-type: none">GaussDB(DWS) cluster

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 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.

Table 3-187 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 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: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.
Synchronization object	<ul style="list-style-type: none"> • Full synchronization supports the synchronization of data, table structures, and 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.

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 database sharding middleware cannot contain the following characters: '<>/\ and non-ASCII characters.• 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 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.

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. ● 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. ● 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. ● 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 incremental synchronization, some DDL statements are supported, including ADD COLUMN, CREATE TABLE, MODIFY COLUMN, CREATE INDEX, DROP INDEX and RENAME INDEX. ● During incremental synchronization, perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. ● During 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 3-176 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: [Dropdown]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown]

* Task Name: [Text Input: DRS-6131]

Description: [Text Area: 0/256]

Table 3-188 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-177 Synchronization instance details

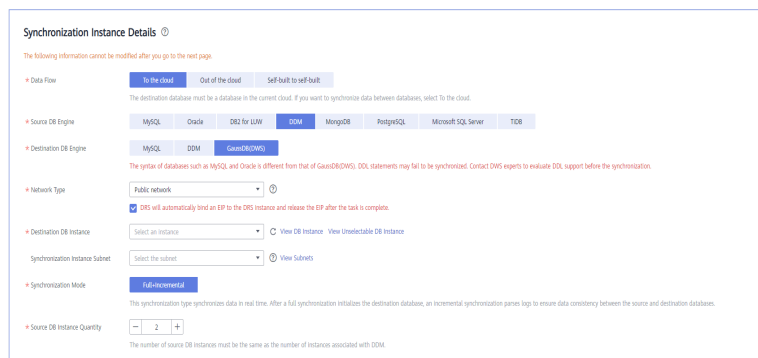


Table 3-189 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Destination DB Instance	An available GaussDB(DWS) instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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
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.

- DRS instance specifications

Figure 3-178 Specifications



Table 3-190 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-179 Enterprise projects and tags

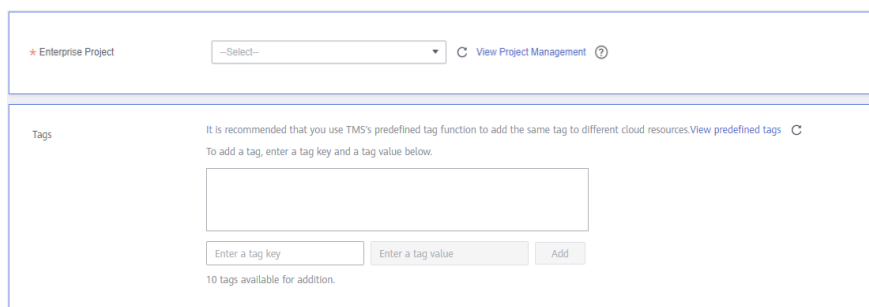


Table 3-191 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-180 Source database information

Source Database

Middleware IP Address or Domain Name

Port

Middleware Username

Database Password

SSL Connection

DB Instance	IP Address or Domain Name ...	Port	Username	Password	SSL Connection
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input style="border: 1px solid gray;" type="password"/>	<input type="checkbox"/>
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input style="border: 1px solid gray;" type="password"/>	<input type="checkbox"/>

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

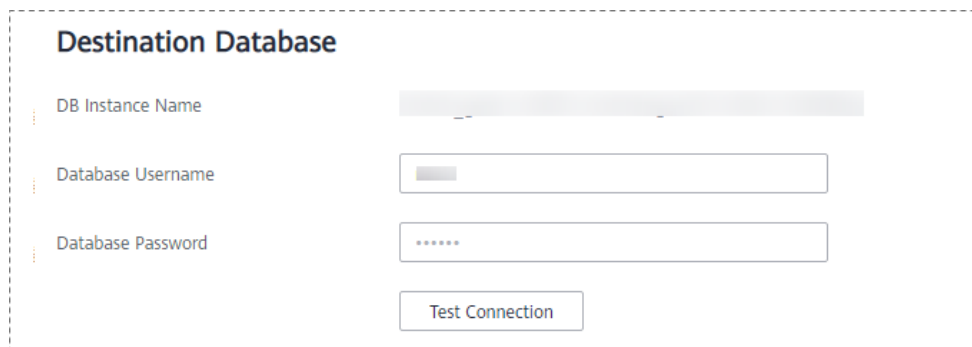
Table 3-192 Source database settings

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 not enabled, 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 3-181 Destination database information



Destination Database

DB Instance Name

Database Username

Database Password

Table 3-193 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.

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.

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 3-182 Synchronization mode

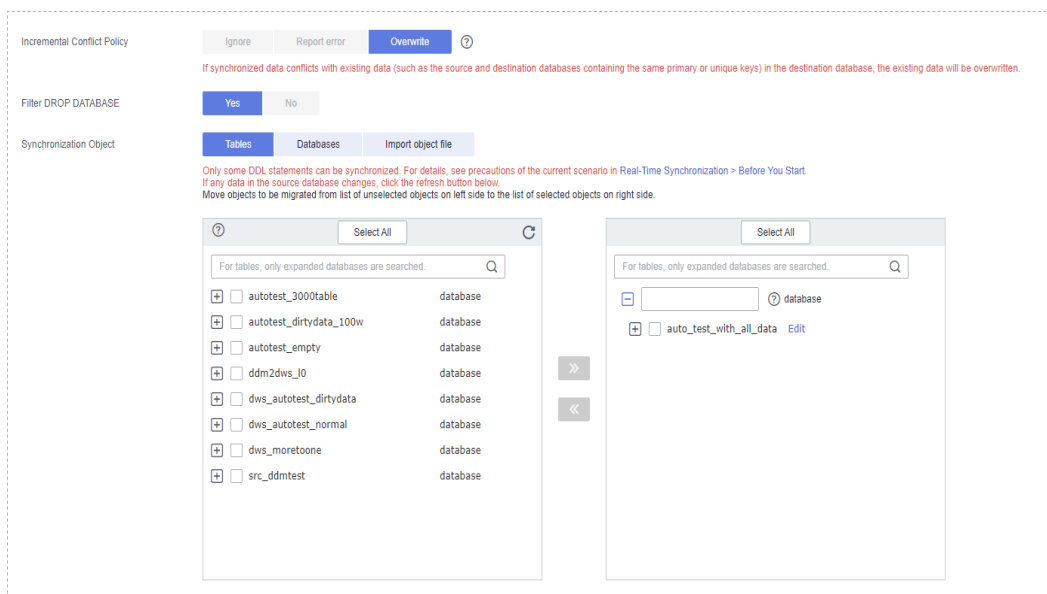




Table 3-194 Synchronization mode and object

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.

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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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](#),

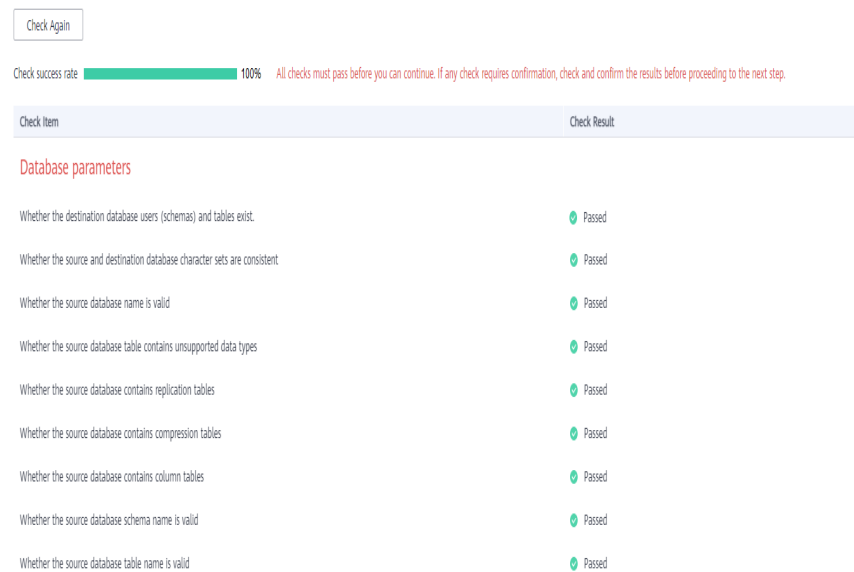
Figure 3-183 Processing data



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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 3-184 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-185 Task startup settings

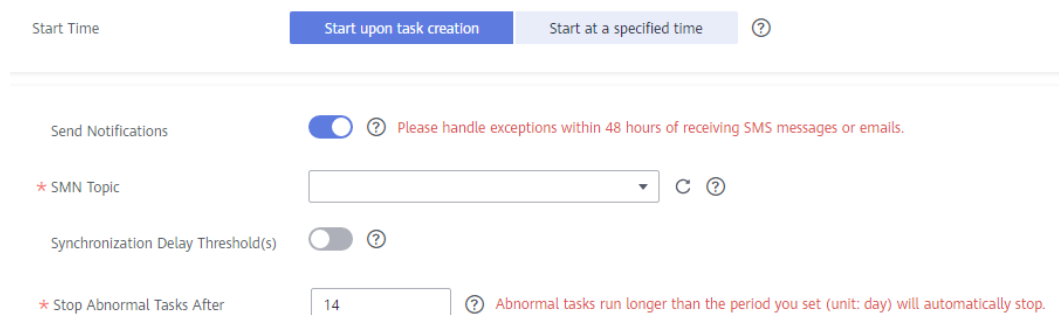



Table 3-195 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.20 From DDM to DDM

Supported Source and Destination Databases

Table 3-196 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">DDM instances	<ul style="list-style-type: none">DDM instances

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 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.

Table 3-197 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 physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The DDM destination database 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 the permission on the database to be synchronized.
Synchronization object	<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 timestamp. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

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 source database binlog 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 the following characters: '<>/\` and non-ASCII characters. • Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. • Enable GTID of the source database.
Destination database	<ul style="list-style-type: none"> • 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. • The destination instance and associated RDS instance must be available. If the RDS instance type is primary/standby, the replication status must be normal. • The associated RDS instance must have sufficient storage space. • The character set of the associated RDS database must be the same as that of the source database. • If the destination instance uses columns of the TIMESTAMP or DATETIME data type as its sharding key, the seconds precision of the column is removed after the synchronization. • The value of AUTO_INCREMENT of a table in the destination database cannot be less than that of AUTO_INCREMENT of a table in the source database.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization 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, 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 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. ● During an incremental synchronization, do not perform the restoration operation on the source database. ● During an incremental synchronization, the following DDL statements are supported: ADD COLUMN, CREATE TABLE, CREATE INDEX, DROP INDEX, RENAME INDEX and ADD INDEX. ● During incremental synchronization, perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. ● During an incremental synchronization of table-level objects, renaming tables is not recommended. ● If the target DDM version is later than 3.0.4.1, DRS automatically updates the start value of the DDM sequence when the task is complete.

Type	Restrictions
	<ul style="list-style-type: none"> Set the expire_log_day parameter for the physical shards of the source middleware to a proper value to ensure that the binlog at the breakpoint does not expire during restoration and that the service can be successfully restored 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 3-186 Synchronization task information

Table 3-198 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-187 Synchronization instance details

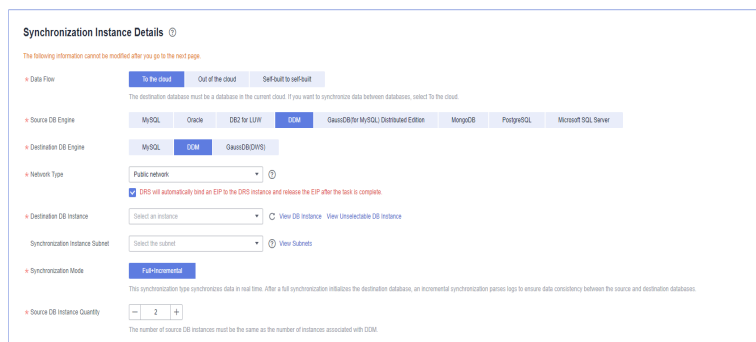


Table 3-199 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DDM .
Destination DB Engine	Select DDM .
Network Type	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Destination DB Instance	The DDM instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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
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.

- DRS instance specifications

Figure 3-188 Specifications



Table 3-200 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 3-189 Enterprise projects and tags

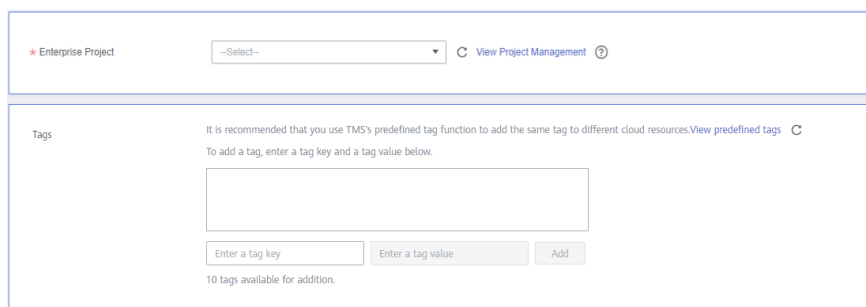


Table 3-201 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-190 Source database information

Source Database



Middleware IP Address or Domain Name

Port

Middleware Username

Database Password 

SSL Connection

DB Instance	IP Address or Domain Name ...	Port	Username	Password	SSL Connection
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input style="border: 1px solid gray; padding: 2px 10px;" type="password"/> 	<input type="checkbox"/>
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input style="border: 1px solid gray; padding: 2px 10px;" type="password"/> 	<input type="checkbox"/>

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

Table 3-202 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 not enabled, 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 3-191 Destination database information

The screenshot shows a configuration window titled "Destination Database". It contains three input fields: "DB Instance Name" (with a dropdown arrow), "Database Username" (with a small grey bar), and "Database Password" (with masked characters "*****"). Below these fields is a "Test Connection" button.

Table 3-203 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.

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.

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

Figure 3-192 Synchronization mode

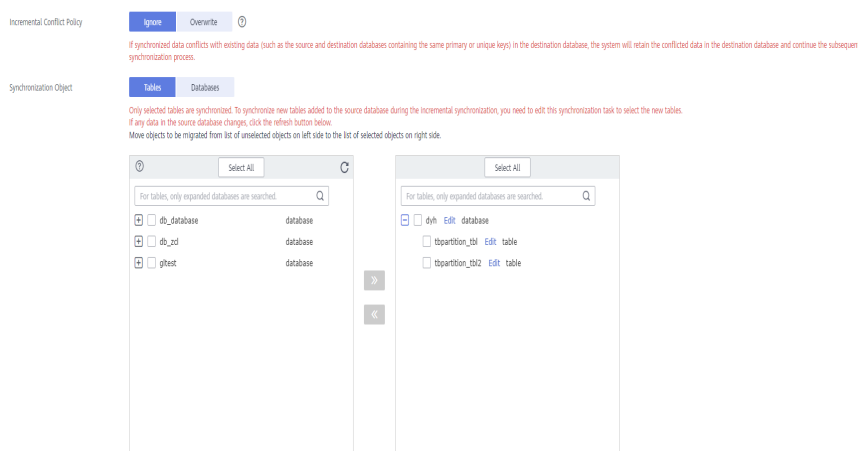



Table 3-204 Synchronization mode and object

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. ● 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

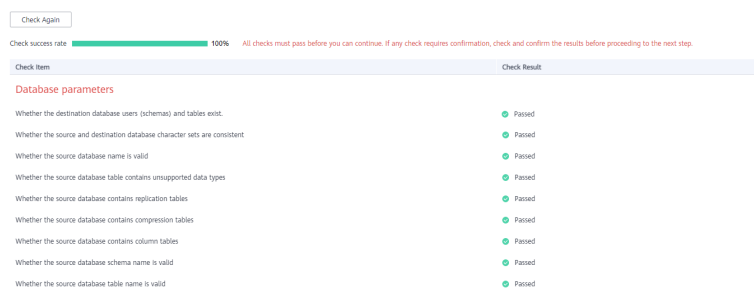
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-193 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-194 Task startup settings





Table 3-205 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.21 From DB2 for LUW to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-206 Supported databases

Source DB	Destination DB
DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB primary/standby

 **NOTE**

Only whitelisted users can use this function.

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 3-207 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. - 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. ● The destination database must have the following permissions: <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. - 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary keys, unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. • User-defined types are not supported. • The XML type is not supported in the incremental phase. • The maximum precision supported by the timestamp type is 6. • Full synchronization and structure migration do not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • Tables whose default values contain expressions of functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • DDL statements executed in the source database cannot be synchronized.
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 on the current cloud. ● 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: <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.

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 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. ● 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 fully synchronized, 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. ● 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.

Type	Restrictions
	<ul style="list-style-type: none"> 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 3-195 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-0131 [Text input]

Description: [Text area]

0/256

Table 3-208 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-196 Synchronization instance details




Table 3-209 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	The public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
Destination DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full and Full+Incremental .

- Enterprise Projects and Tags

Figure 3-197 Enterprise projects and tags




Table 3-210 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-198 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

Table 3-211 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.

 **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 3-199 Destination database information

The screenshot shows a form titled "Destination Database" enclosed in a dashed border. It contains three input fields: "DB Instance Name" (with a greyed-out value), "Database Username" (with a small greyed-out value), and "Database Password" (with six asterisks). Below these fields is a "Test Connection" button.

Table 3-212 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 3-200 Synchronization mode

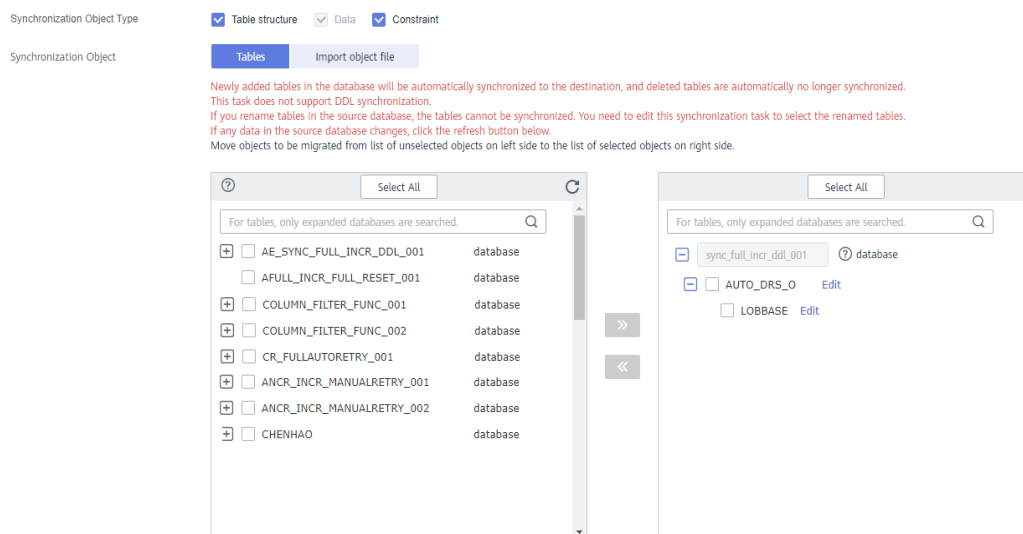



Table 3-213 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>Available options: Table structure, Data, and Constraint</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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

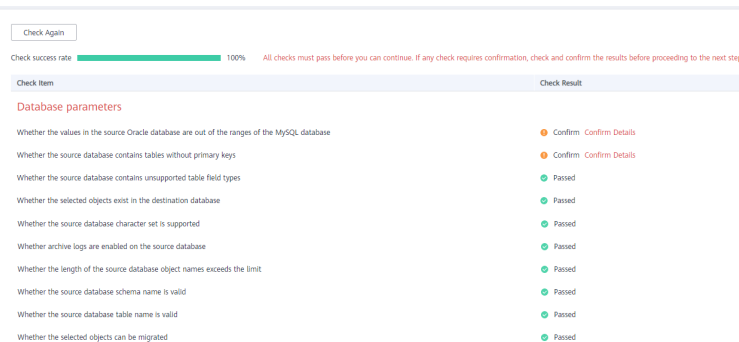
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-201 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-202 Task startup settings





Table 3-214 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.22 From DB2 for LUW to GaussDB Distributed

Supported Source and Destination Databases

Table 3-215 Supported databases

Source DB	Destination DB
DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB distributed

NOTE

Only whitelisted users can use this function.

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 3-216 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. - 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. ● The destination database must have the following permissions: <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. - 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● User-defined types are not supported. ● The XML type is not supported in the incremental phase. ● The maximum precision supported by the timestamp type is 6. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● Tables whose default values contain expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized.
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 on the current cloud.● 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:<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.

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 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.• 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 fully synchronized, 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.• 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.

Type	Restrictions
	<ul style="list-style-type: none"> 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 3-203 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: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 [Text input]

Description: [Text area]

0/256

Table 3-217 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-204 Synchronization instance details

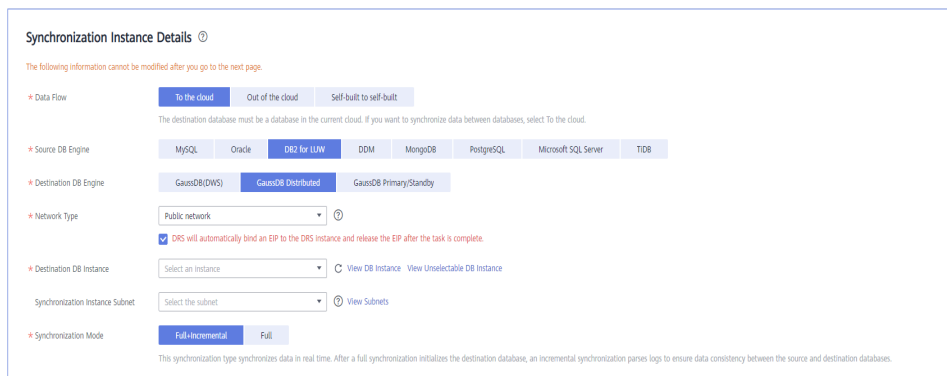


Table 3-218 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select DB2 for LUW.
Destination DB Engine	Select GaussDB Distributed .
Network Type	The public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
Destination DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full and Full+Incremental .

- Enterprise Projects and Tags

Figure 3-205 Enterprise projects and tags




Table 3-219 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

- 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 3-206 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

Table 3-220 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.

 **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 3-207 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields and a button. The fields are: "DB Instance Name" (a long text field), "Database Username" (a text field), and "Database Password" (a password field with masked characters). Below the fields is a "Test Connection" button.

Table 3-221 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 3-208 Synchronization mode

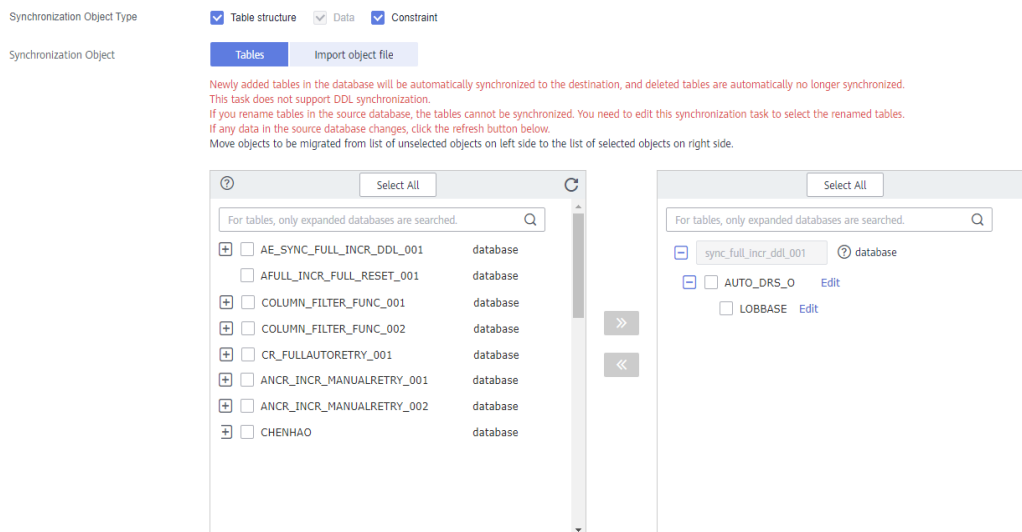



Table 3-222 Synchronization Object

Parameter	Description
Synchronization Object Type	<p>Available options: Table structure, Data, and Constraint</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.

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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

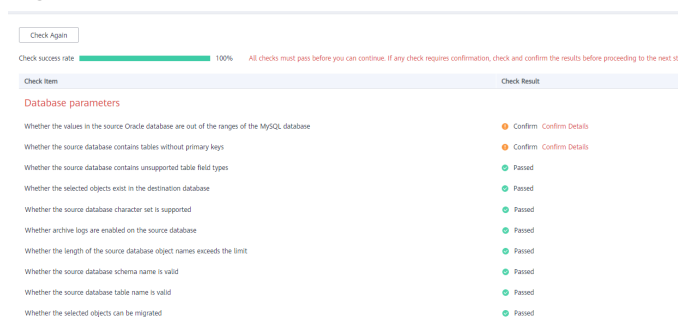
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-209 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-210 Task startup settings





Table 3-223 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.23 From DB2 for LUW to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-224 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5 	<ul style="list-style-type: none"> • GaussDB(DWS)

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 3-225 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. - 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: <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.

Type	Restrictions
Synchronization object constraints	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, unique constraints, and data can be synchronized. Tables can be synchronized in real time during incremental synchronization. ● User-defined types are not supported. ● The XML type is not supported in the incremental phase. ● The maximum precision supported by the timestamp type is 6. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● Tables whose default values contain expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● 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.
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. ● If you select full+incremental synchronization, ensure that archive log of the source database is enabled. ● Only the GBK and UTF8 character sets can be synchronized.

Type	Restrictions
Destination database	<ul style="list-style-type: none">● 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:<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.

Type	Restrictions
Precautions	<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 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.• 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 fully synchronized, 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.• 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.

Type	Restrictions
	<ul style="list-style-type: none"> 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 3-211 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. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Project: [Dropdown menu]

* Task Name: [Text input: DRS-0131] ⓘ

Description: [Text area: 0/256] ⓘ

Table 3-226 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-212 Synchronization instance information

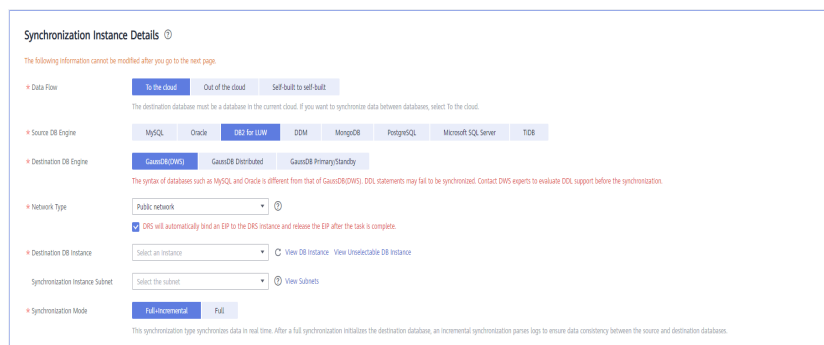


Table 3-227 Synchronization instance settings

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	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Destination DB Instance	The GaussDB(DWS) instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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>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. <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.

- Enterprise Projects and Tags

Figure 3-213 Enterprise projects and tags

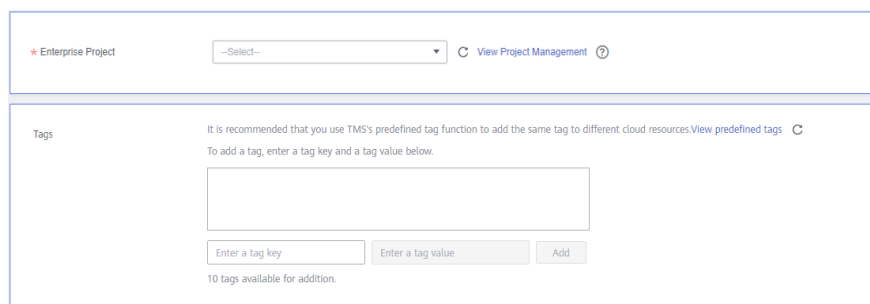


Table 3-228 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-214 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

Table 3-229 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.

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. 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.

 **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 3-215 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 


 Test successful

Table 3-230 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 3-216 Synchronization Mode

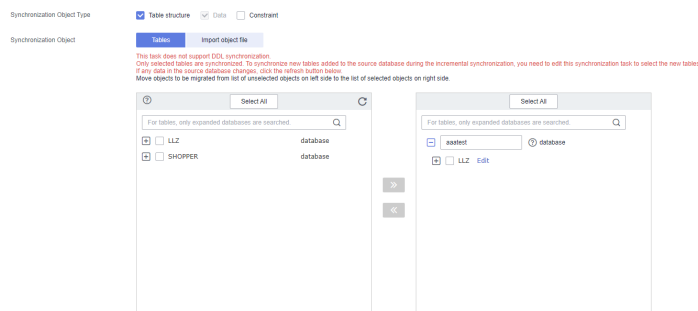



Table 3-231 Synchronization object

Parameter	Description
Synchronization Object Type	<p>Available options: Table structure, Data, and Constraint</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.

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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

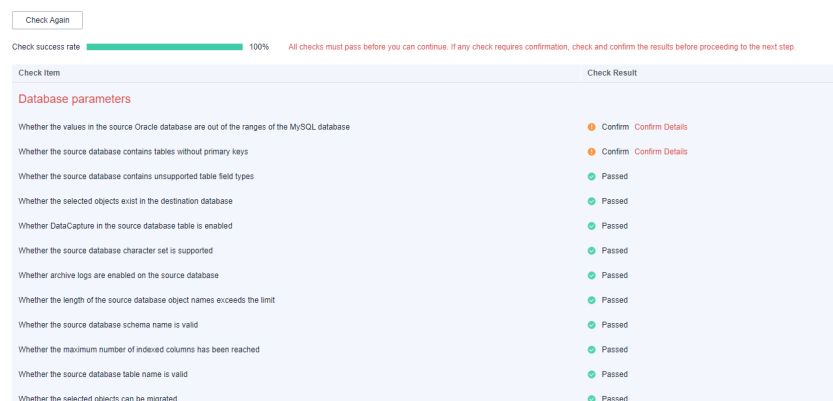
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-217 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-218 Task startup settings

Table 3-232 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.24 From TiDB to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-233 Supported databases

Source DB	Destination DB
TiDB 4.0.0 and later (excluding the development version)	GaussDB(for MySQL)

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 3-234](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-234 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none">• 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.• Scope of full synchronization<ul style="list-style-type: none">- The structure, sequence, data, constraints, and indexes of selected tables.- The names of the databases and tables to be synchronized cannot contain non-ASCII characters or special characters '< > \- The database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1.- 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.• Scope of incremental synchronization<ul style="list-style-type: none">- Some DML statements, including INSERT, UPDATE, and DELETE- Some DDL statements, including CREATE TABLE, DROP TABLE, ALTER TABLE, RENAME TABLE and TRUNCATE TABLE

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 3-235](#). 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.

Table 3-235 Database account permission

Type	Full+Incremental
Source database user	SELECT and CONFIG permissions.
Destination database user	SELECT, CREATE, DROP, INSERT, DELETE, UPDATE, ALTER, REFERENCE and INDEX permissions.

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 3-236 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">– 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● CAUTION<ul style="list-style-type: none">● TiDB Binlog is incompatible with the following features in TiDB 5.0: 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: TiDB system variable tidb_enable_amend_pessimistic_txn: After TiDB Binlog is enabled, the data replicated from the TiDB Binlog may be inconsistent.● Source database object requirements:<ul style="list-style-type: none">– The names of the source objects to be synchronized cannot contain non-ASCII characters or special characters '<' '>' '\'● Destination database parameter requirements:<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.● Destination database object requirements:<ul style="list-style-type: none">– 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: 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"> - Triggers enabled in the destination database cannot be associated with synchronization tables. <p>● Other notes:</p> <ul style="list-style-type: none"> - You can add additional objects during an incremental synchronization.
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.
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.
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.

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

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 3-219 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-0131 [Text input]

Description: [Text area]

0/256

Table 3-237 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-220 Synchronization instance details




Table 3-238 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select TiDB .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	Available options: Public network and VPN or Direct Connect . Public network is used as an example.
Destination 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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 3-221 Enterprise projects and tags

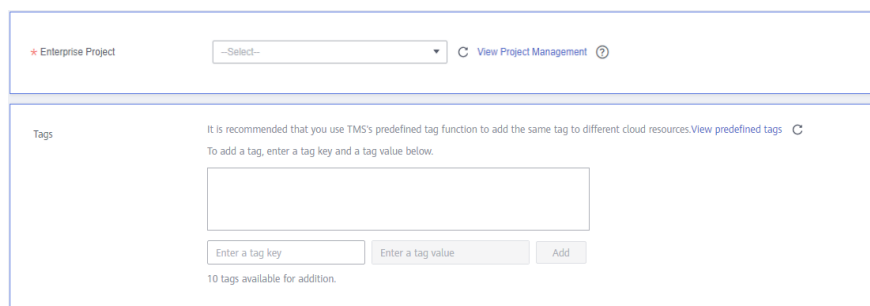


Table 3-239 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-222 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 3-240 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 not enabled, your data may be at risk.

Figure 3-223 Kafka information

Kafka Information

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

● Test successful

Table 3-241 Kafka settings


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.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Figure 3-224 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 3-242 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 3-225 Synchronization Mode

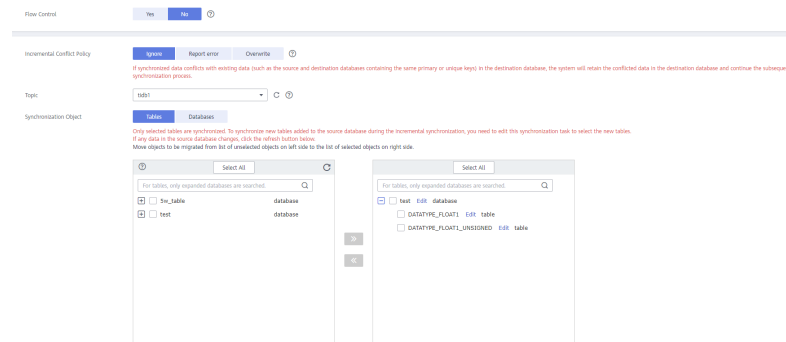
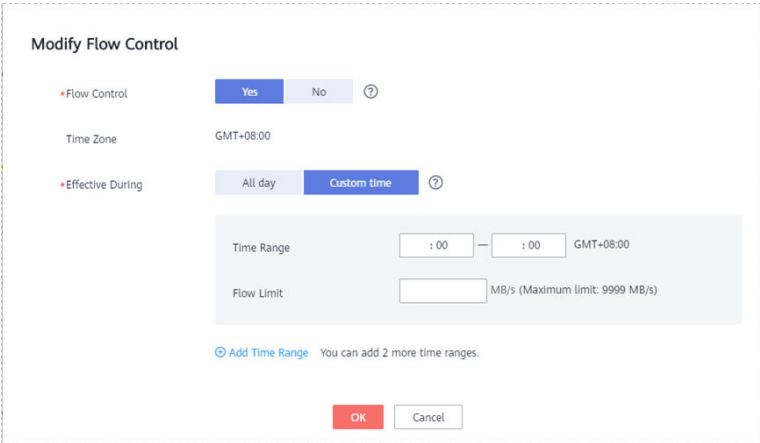



Table 3-243 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-226 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.
Topic	Select a topic that stores TiDB binlogs in the Kafka, or synchronization will fail.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

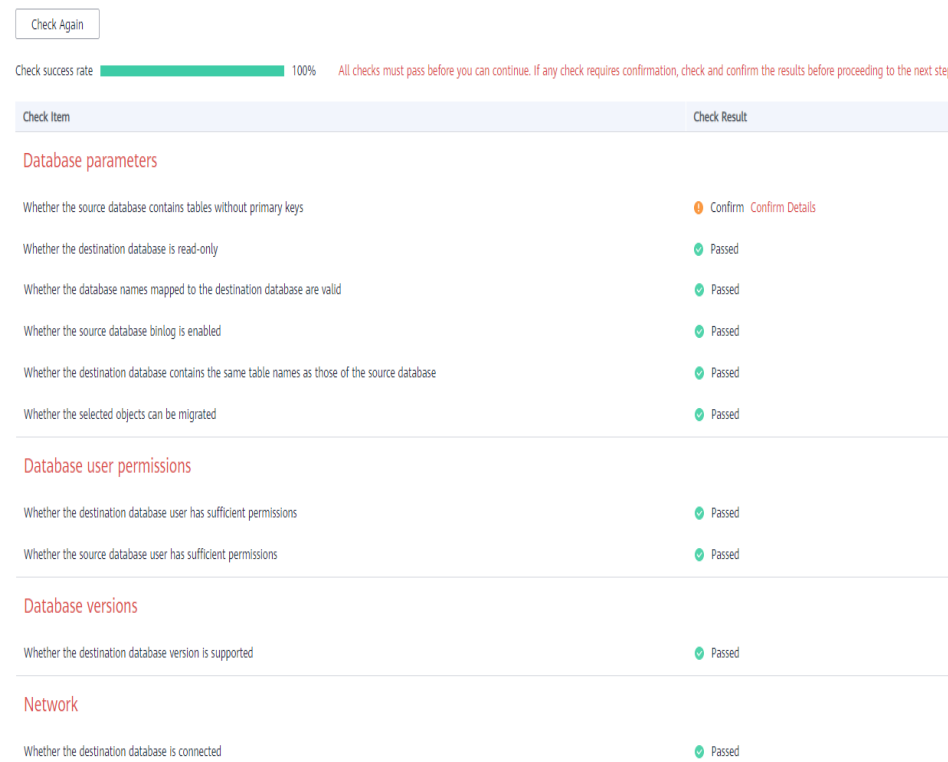
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-227 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-228 Task startup settings





Table 3-244 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.25 From Microsoft SQL Server to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-245 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on an ECS (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on other clouds (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) RDS for SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) 	<ul style="list-style-type: none"> GaussDB(DWS) cluster

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-246 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-246 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">- Structures, data, and indexes of a table are supported.- 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.• 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 3-247](#). 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.

Table 3-247 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.

Suggestions

 **CAUTION**

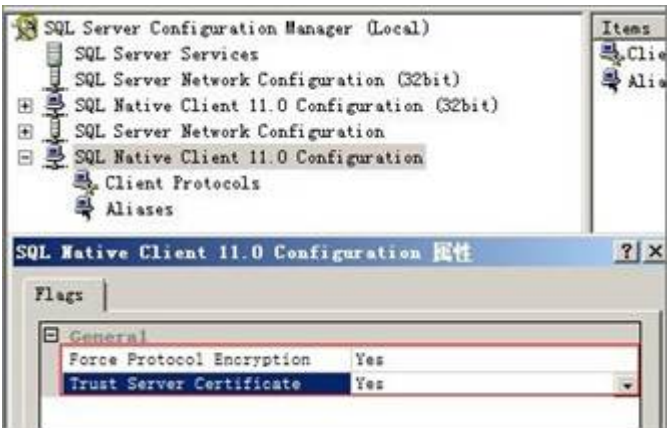
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- 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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 3-248 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. - 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 3-229. <p style="text-align: center;">Figure 3-229 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. - The database name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The table name in the source database cannot contain special characters '<>' ● 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 SQL Server system databases). - 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

Type	Constraints
	<p>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"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - 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.
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.
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. ● 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.

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.

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

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 3-230 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: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-B131 [Help icon]

Description: [Text area] 0/256 [Help icon]

Table 3-249 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-231 Synchronization instance details

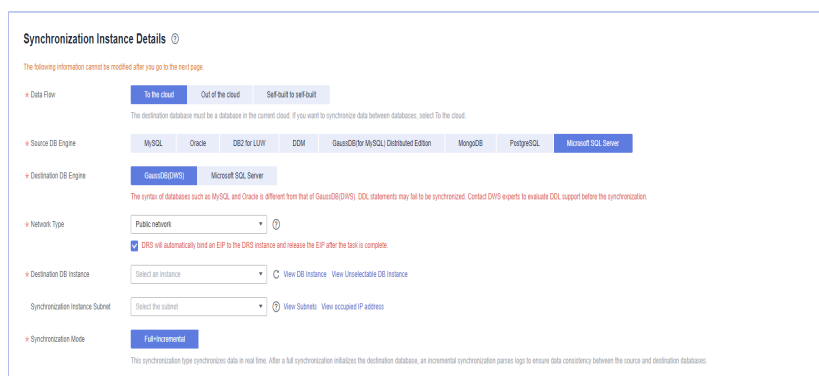


Table 3-250 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example.
Destination DB Instance	The destination is a GaussDB(DWS) instance.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 3-232 Enterprise projects and tags

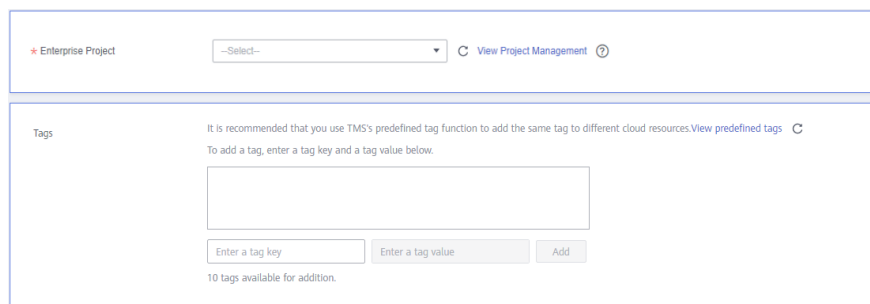


Table 3-251 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-233 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 3-252 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

Parameter	Description
Database Username	The username for logging in to the source Microsoft SQL Server 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 not enabled, your data may be at risk.

Figure 3-234 Destination database information

Table 3-253 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.
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 3-235 Synchronization mode

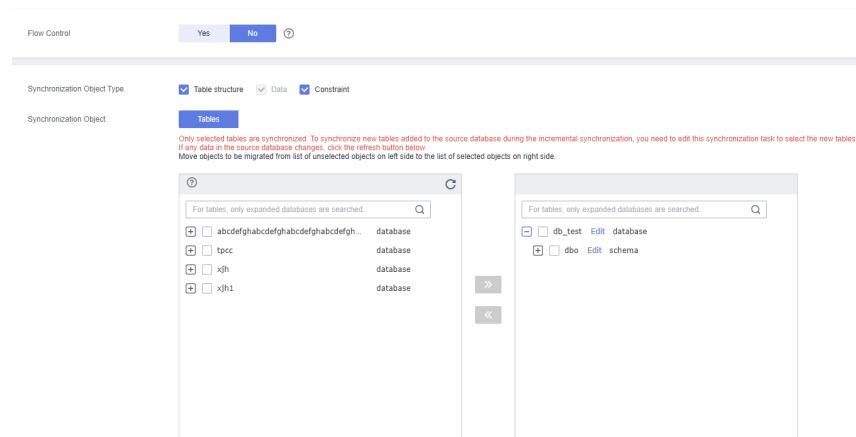
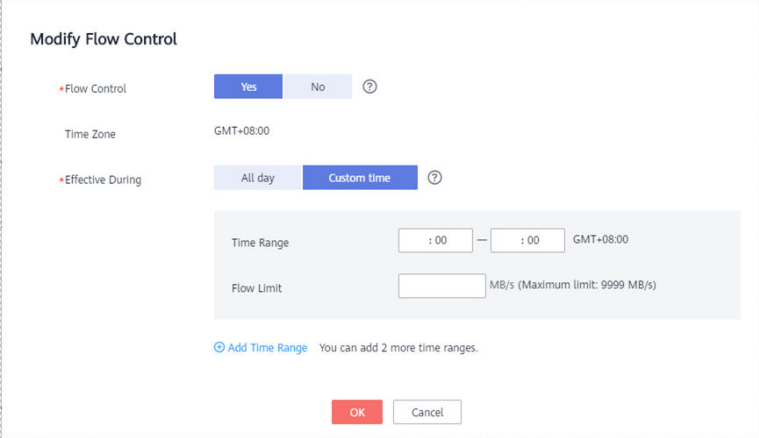



Table 3-254 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-236 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>Available options: Table structure, Data, and Constraint</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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

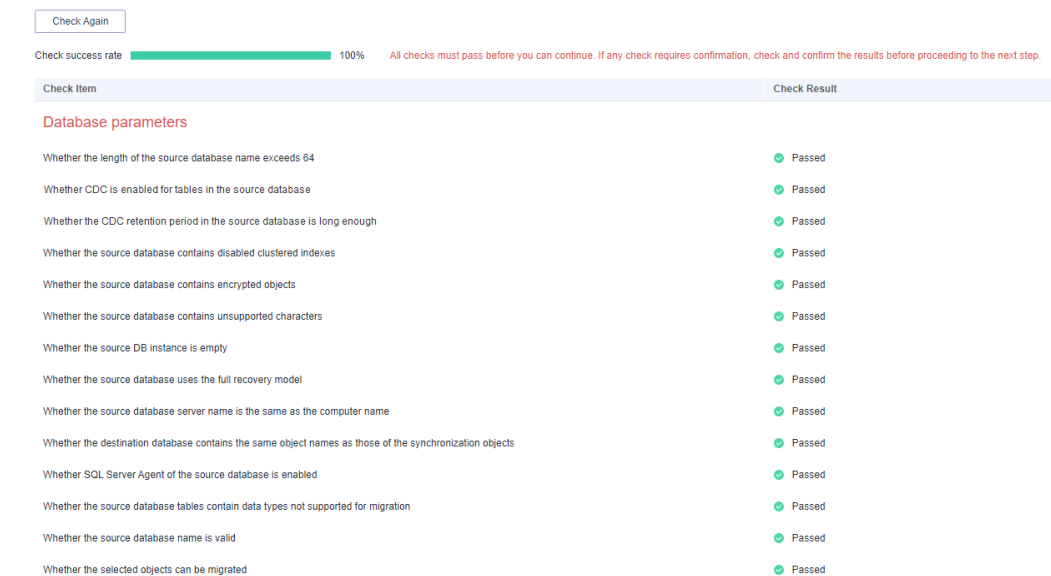
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-237 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-238 Task startup settings

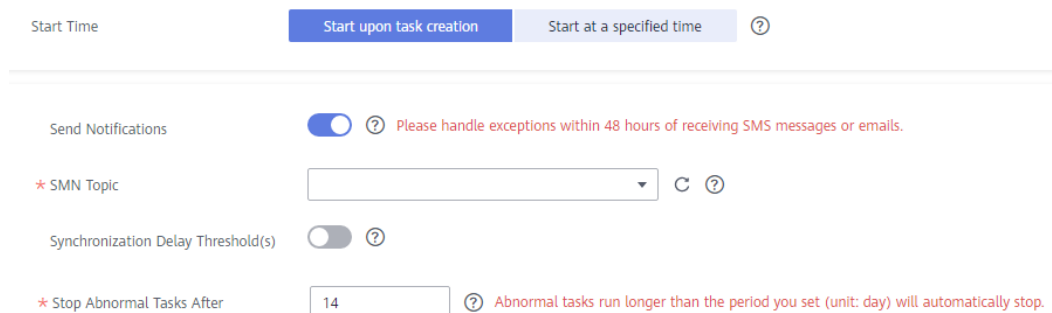



Table 3-255 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.26 From Microsoft SQL Server to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-256 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on an ECS (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on other clouds (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) RDS for SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) 	GaussDB primary/standby

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-257 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-257 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">- Structures, data, and indexes of a table are supported.- 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.• 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 3-258](#). 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.

Table 3-258 Database account permission

Type	Full+Incremental
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"> • 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. • 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.

Suggestions

 **CAUTION**

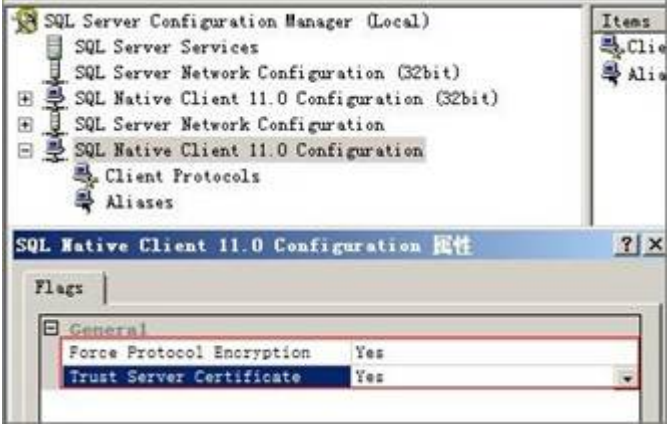
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- 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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 3-259 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. - 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 3-239. <p style="text-align: center;">Figure 3-239 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. - The database name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The table name in the source database cannot contain special characters '<>' ● 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 SQL Server 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. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the

Type	Constraints
	<p>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.</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.
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. • 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.

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

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 3-240 Synchronization task information

Table 3-260 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-241 Synchronization instance details

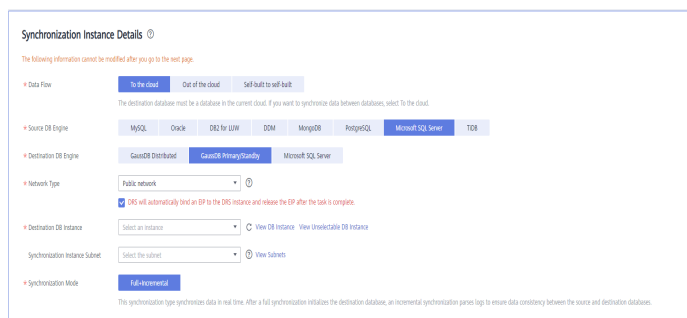


Table 3-261 Synchronization instance settings

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.
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 3-242 Enterprise projects and tags

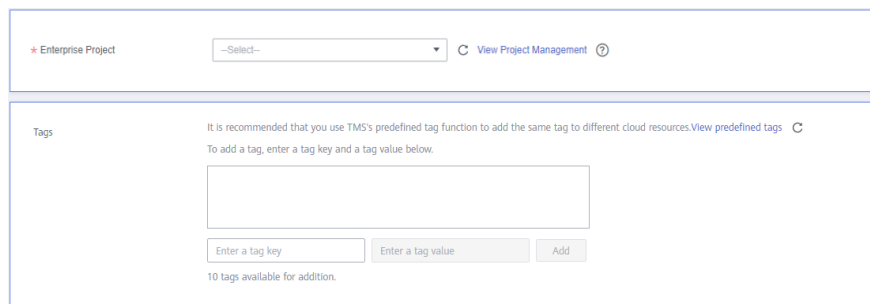


Table 3-262 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-243 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 3-263 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.
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 not enabled, your data may be at risk.

Figure 3-244 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-264 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 3-245 Synchronization mode


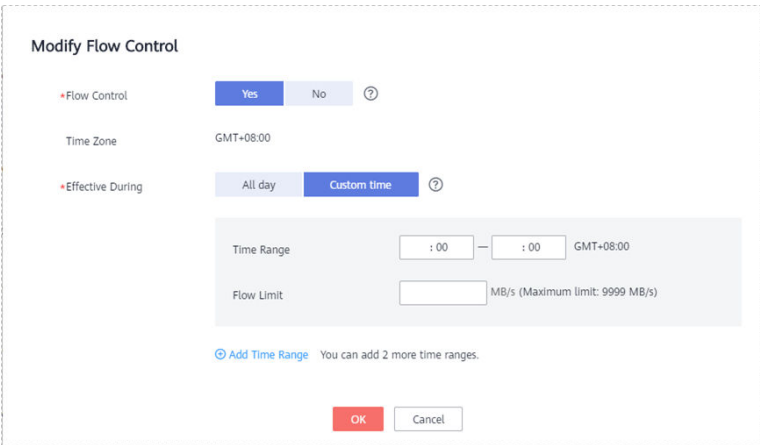



Table 3-265 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-246 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
Synchronization Object Type	<p>Available options: Table structure, Data, and Constraint</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. • 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

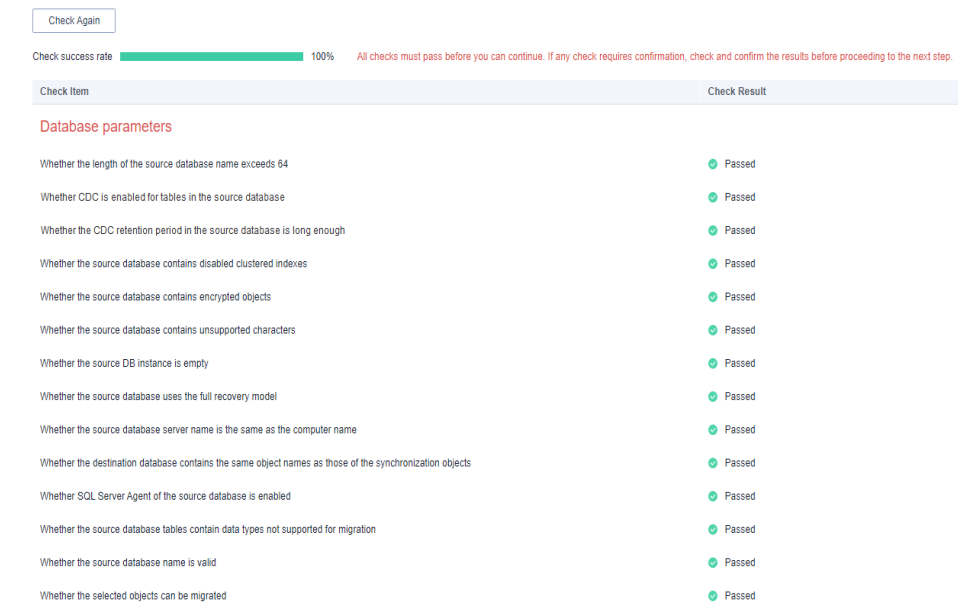
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-247 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-248 Task startup settings

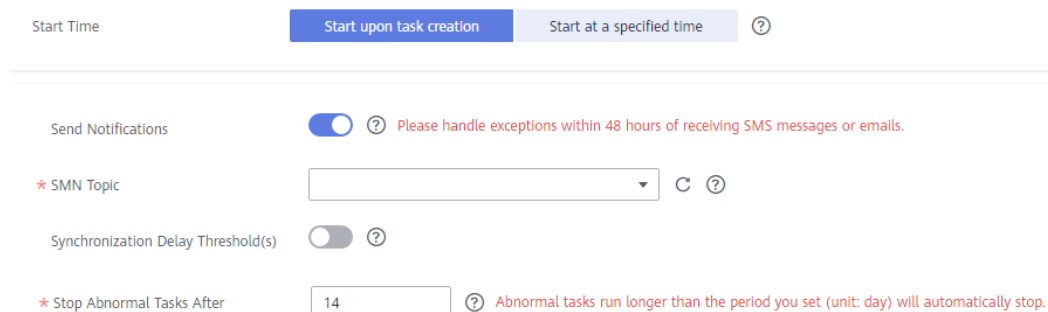



Table 3-266 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.27 From Microsoft SQL Server to GaussDB Distributed

Supported Source and Destination Databases

Table 3-267 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on an ECS (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on other clouds (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) RDS for SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) 	GaussDB distributed

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-268 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-268 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">- Structures, data, and indexes of a table are supported.- 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.● 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 3-269](#). 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.

Table 3-269 Database account permission

Type	Full+Incremental
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">• 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.• 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.

Suggestions

⚠ CAUTION

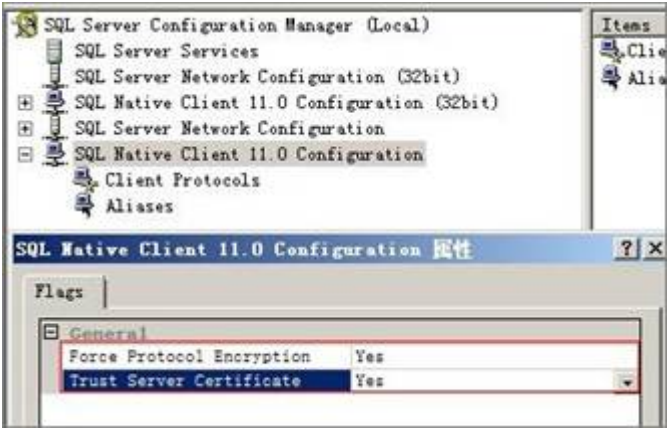
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- 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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 3-270 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. - 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 3-249. <p style="text-align: center;">Figure 3-249 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. - The database name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The table name in the source database cannot contain special characters '<>' ● 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 SQL Server 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. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the

Type	Constraints
	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.
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.
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. 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.

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

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 3-250 Synchronization task information

Table 3-271 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-251 Synchronization instance details

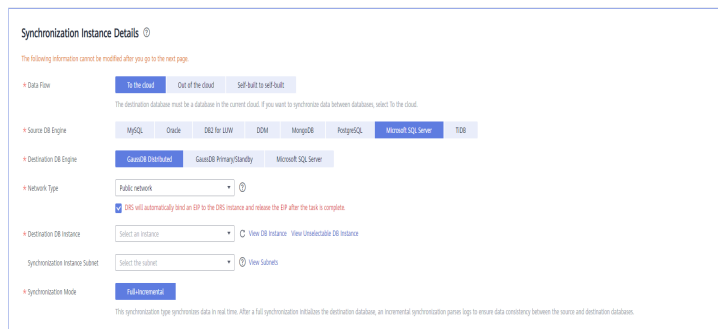


Table 3-272 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select GaussDB Distributed .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example.
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 3-252 Enterprise projects and tags




Table 3-273 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-253 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 3-274 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.
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 not enabled, your data may be at risk.

Figure 3-254 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-275 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 3-255 Synchronization mode

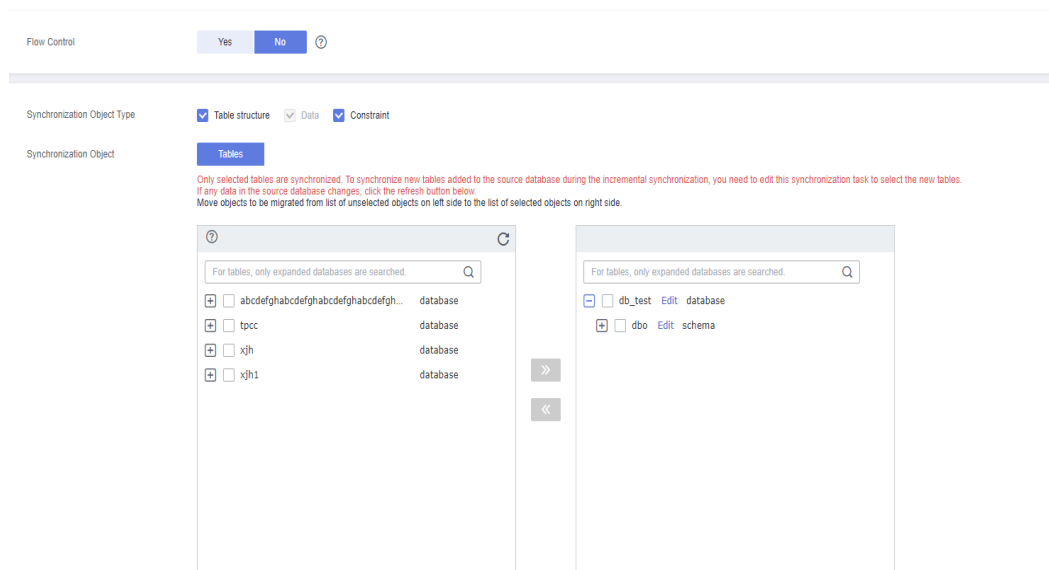
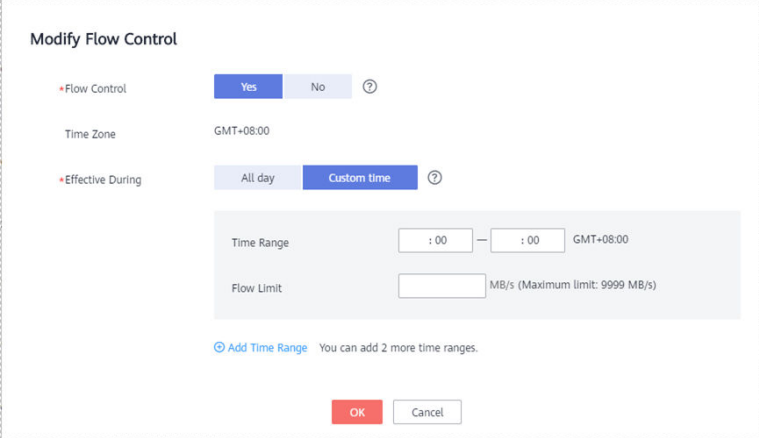



Table 3-276 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-256 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>Available options: Table structure, Data, and Constraint</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. • 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

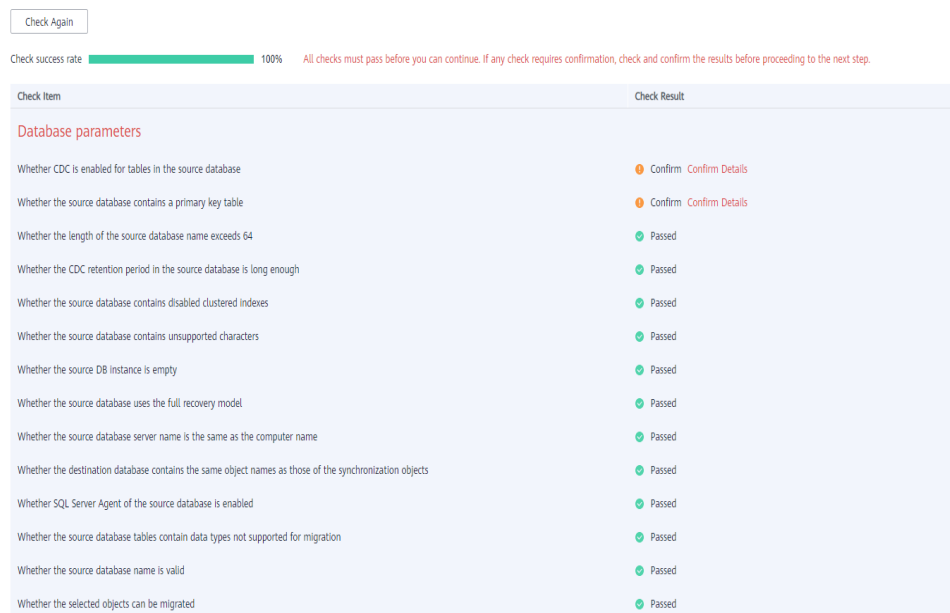
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-257 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-258 Task startup settings

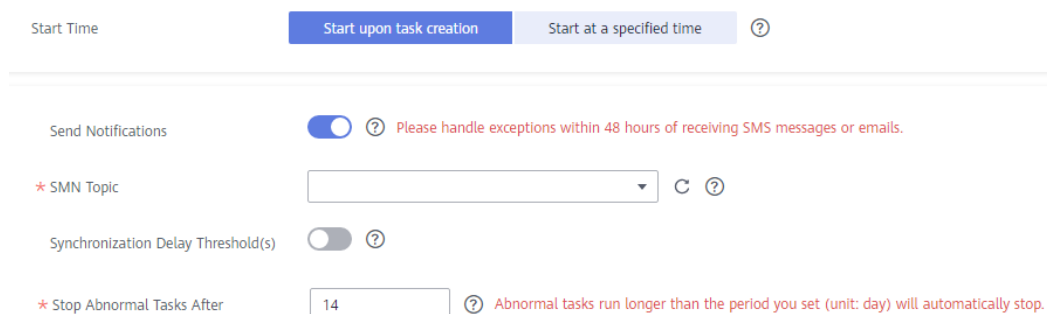



Table 3-277 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.28 From Microsoft SQL Server to Microsoft SQL Server

Supported Source and Destination Databases

Table 3-278 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on an ECS (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) Microsoft SQL Server database on other clouds (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) RDS for SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) 	<ul style="list-style-type: none"> RDS for SQL Server (Enterprise Edition 2012 or later, Standard Edition 2016 SP1 or later) <p>NOTE The major version of the destination database must be the same as or later than that of the source database.</p>

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 3-279 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-279 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"> - Structures, data, and indexes of a table are supported. - 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. ● 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 3-280](#). 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.

Table 3-280 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	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.

Suggestions

CAUTION

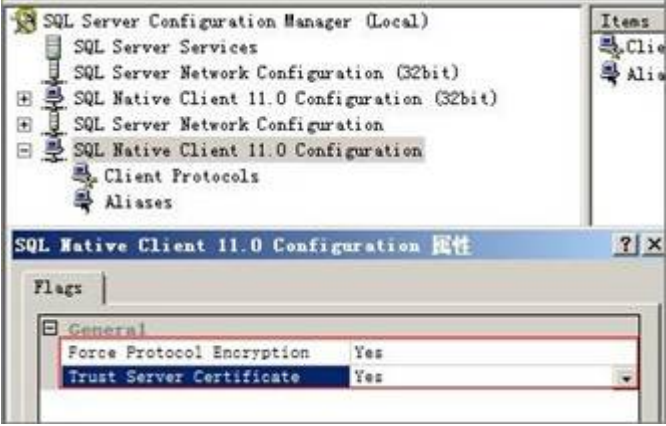
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- 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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 3-281 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. - 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 3-259. <p style="text-align: center;">Figure 3-259 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 database name can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - The table name in the source database cannot contain special characters '<>' ● Destination database object requirements: <ul style="list-style-type: none"> - The destination DB instance must have sufficient storage space. ● Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - 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.

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.
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. 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.

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

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 3-260 Synchronization task information

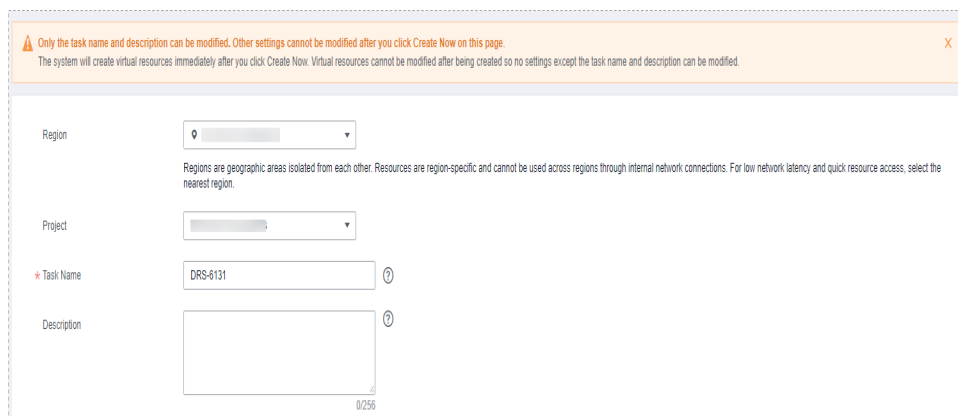


Table 3-282 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-261 Synchronization instance details

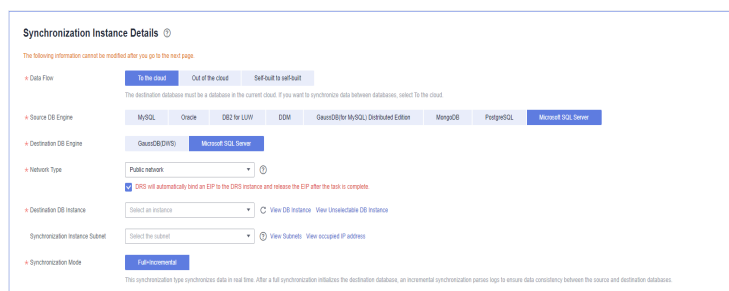


Table 3-283 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Microsoft SQL Server .
Destination DB Engine	Select Microsoft SQL Server .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example.
Destination DB Instance	Select an RDS for SQL Server instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- Enterprise Projects and Tags

Figure 3-262 Enterprise projects and tags




Table 3-284 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">- This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- 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 ends.

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 3-263 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 3-285 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.
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 not enabled, your data may be at risk.

Figure 3-264 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-286 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 3-265 Synchronization Mode

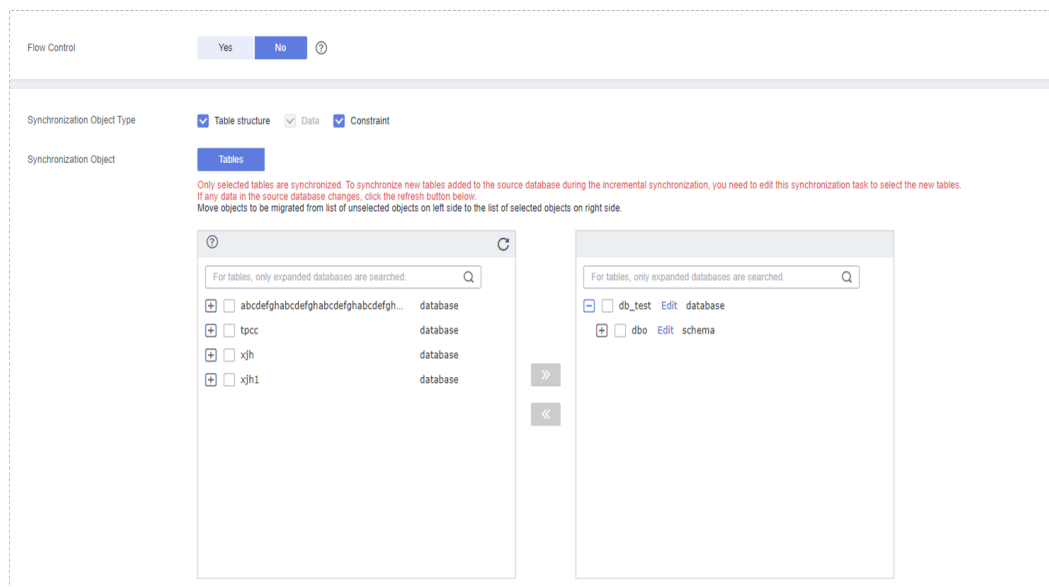
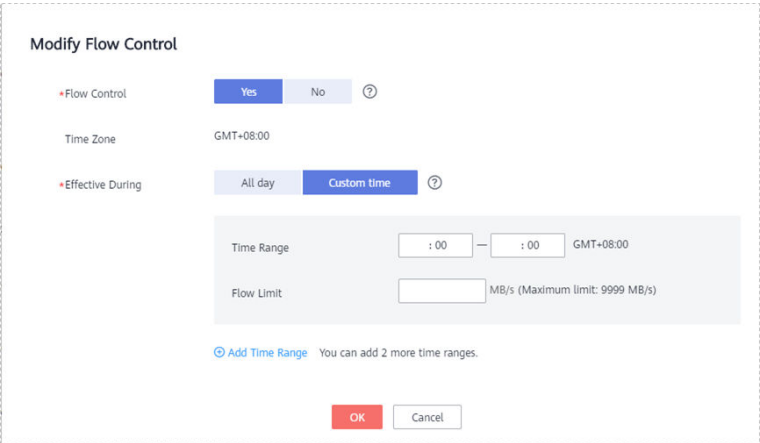



Table 3-287 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-266 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>Available options: Table structure, Data, and Constraint</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. • 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

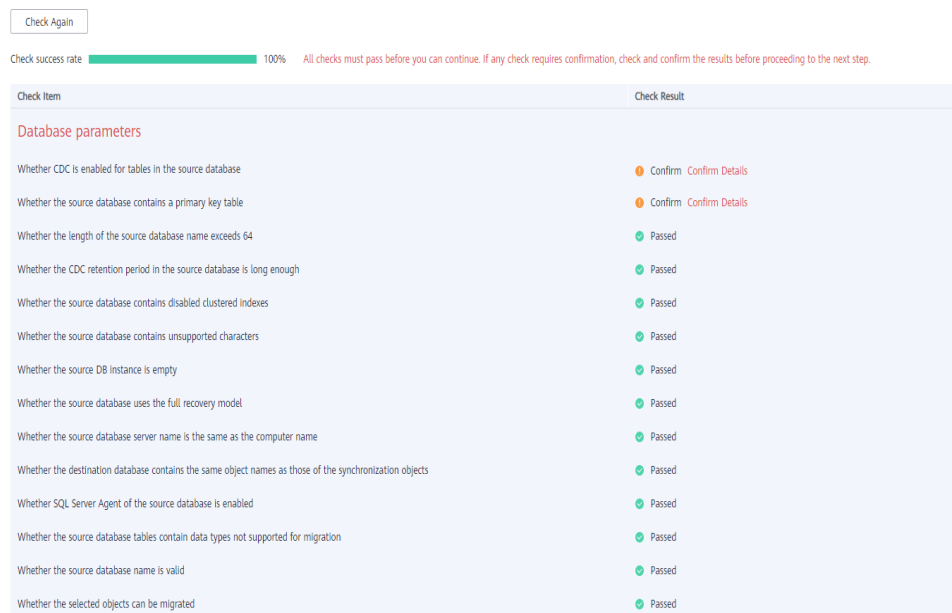
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 3-267 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-268 Task startup settings





Table 3-288 Task startup settings

Parameter	Description
Started 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>SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.</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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

3.29 From MongoDB to DDS

Supported Source and Destination Databases

Table 3-289 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises MongoDB databases• MongoDB databases on an ECS• MongoDB database on other clouds• DDS instances	<ul style="list-style-type: none">• DDS instances

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](#).

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.
- During the synchronization, ensure that no data is written to the destination database to keep data consistency before and after the synchronization.
- Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 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.

Table 3-290 Precautions

Type	Restrictions
Database permissions	<p>Source database (minimum permissions):</p> <ul style="list-style-type: none"> • Permission requirements for full plus 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. <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.</p>
Synchronization object constraints	<ul style="list-style-type: none"> • Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. • The system database (such as local, admin, and config) and system collection cannot be synchronized. If service data is in the 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. • The mapped synchronization object name can contain 1 to 63 characters. The following characters are not allowed <code>\/."\$<></code>

Type	Restrictions
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 or view name cannot start with system. or contain the dollar sign (\$).• The source cannot be a GaussDB(for Mongo) instance.
Destination database	<ul style="list-style-type: none">• The destination DB instance is running properly.• The destination DB instance must have sufficient storage space.• Data cannot be synchronized from a newer version database to an older version database.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: collections referenced by views, and views referenced by views ● Replica set: The MongoDB replica set instance must be available and have primary nodes. ● 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 - The <code>convertToCapped</code>, <code>collMod</code>, and <code>renameCollection</code> commands are supported. ● 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 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 ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization.

Type	Restrictions
	<ul style="list-style-type: none"> • 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 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. • 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 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 3-269 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-6131

Description: [Text area]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Table 3-291 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 3-270 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

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 | MariaDB | **MongoDB** | PostgreSQL | Microsoft SQL Server | TiDB

* Destination DB Engine: **DRS**

* Network Type: VPC

* Destination DB Instance: No DB Instance available | View DB Instance | View Unselectable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnets

* 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.

* Source DB Instance Type: **Non-cluster**

Table 3-292 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 MongoDB .
Destination DB Engine	Select DDS .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network .
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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	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. 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.
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.

- Enterprise Projects and Tags

Figure 3-271 Enterprise projects and tags




Table 3-293 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 3-272 Source database information

Source Database

Database Type: Self-built on ECS **DDS DB Instance**

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

Authentication Database:

Database Username:

Database Password:

🟢 Test successful

Table 3-294 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.

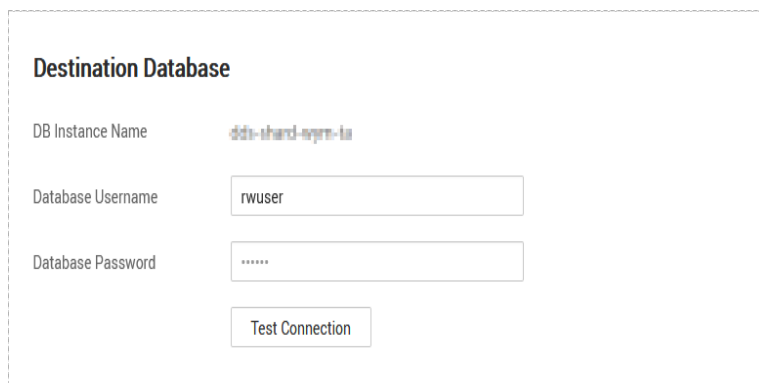
Parameter	Description
SSL Connection	<p>Whether to encrypt the database connections. This option is available if the source database type is Self-built on ECS. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is not enabled, 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 3-273 Destination database information



Destination Database

DB Instance Name: `ddt-shard-svr-1a`

Database Username:

Database Password:

Table 3-295 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.
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 3-274 Synchronization objects

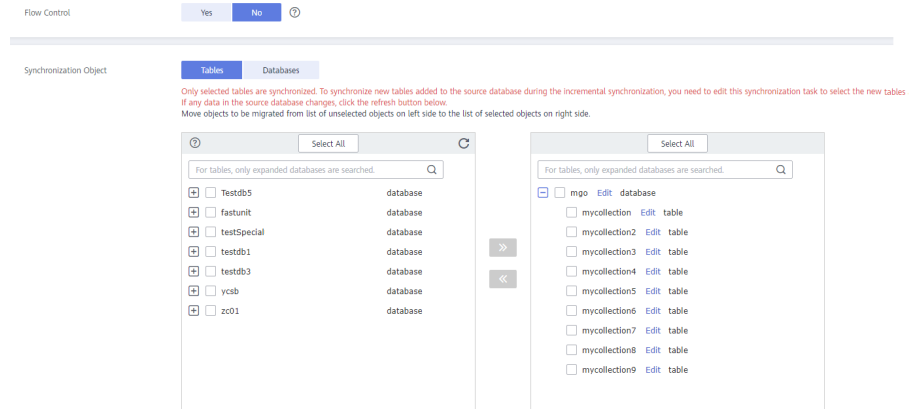
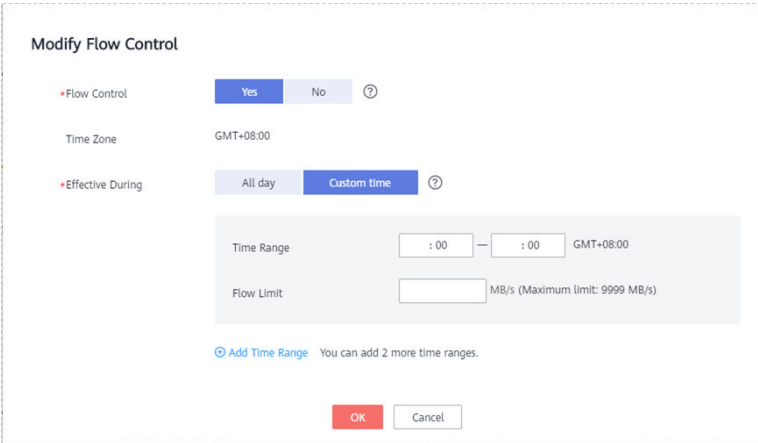




Table 3-296 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 3-275 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 Mapping Object Names. The mapped name can contain 1 to 63 characters. The following characters are not allowed <code>\. "\$ <></code></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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

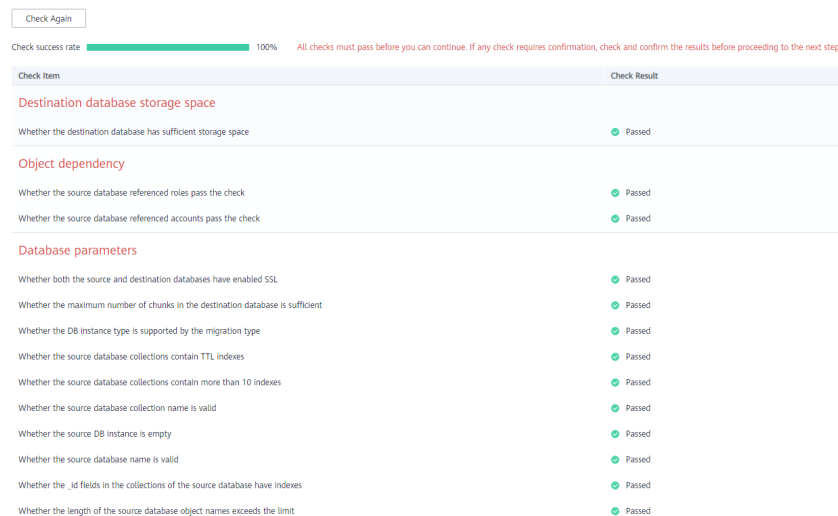
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) 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.

Figure 3-276 Pre-check



Step 6 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 3-277 Task startup settings

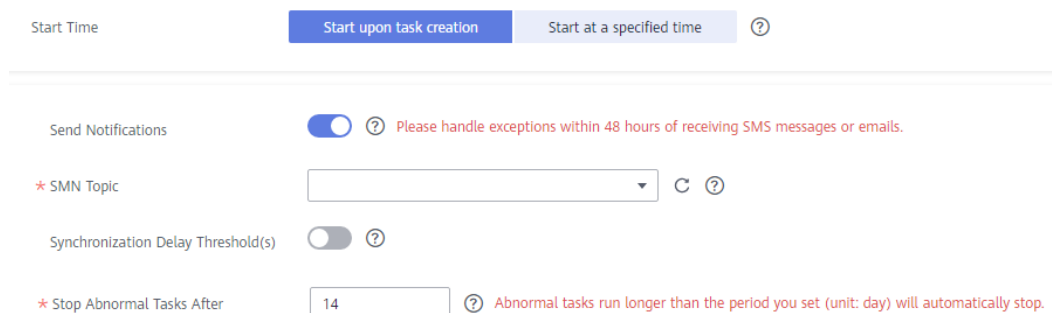



Table 3-297 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----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	<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECSMySQL databases on other cloudsRDS for MySQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 on source databases. 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 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.

Table 4-2 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: 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, stored 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, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects. ● 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. ● Tables with storage engine different to MyISAM and InnoDB 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 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.

Type	Restrictions
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 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. ● 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.

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. ● 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 ● Only MySQL to MySQL synchronization supports many-to-one synchronization. During table-level many-to-one synchronization, tables without primary keys cannot exist in the source database. ● 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 and destination DB instances are RDS for MySQL instances, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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. ● 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 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: <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. ● Data inconsistency may occur when the MyISAM table is modified during a full 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.

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

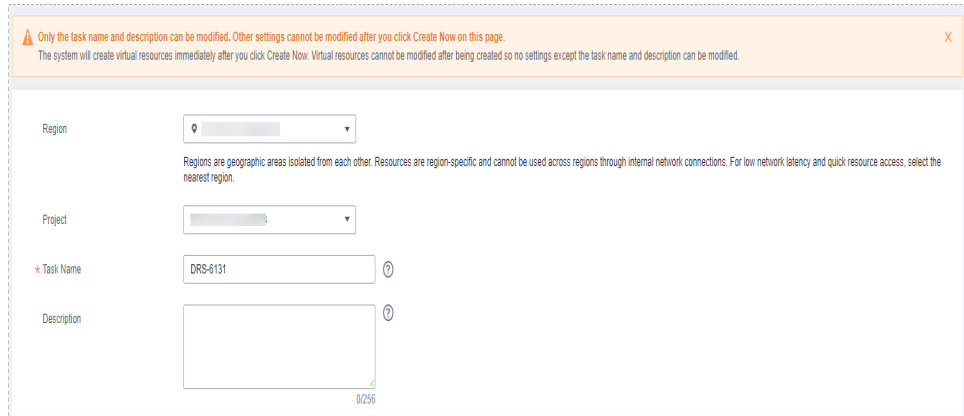


Table 4-3 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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

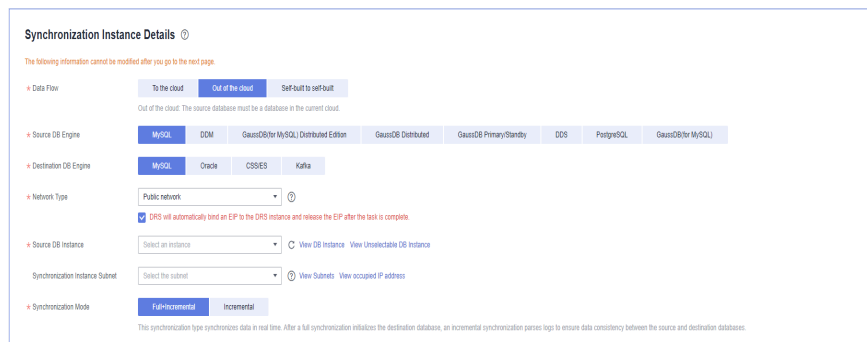


Table 4-4 Synchronization instance settings

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	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source 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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full+Incremental and Incremental <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.

- DRS instance specifications

Figure 4-3 Specifications



Table 4-5 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-4 Enterprise projects and tags

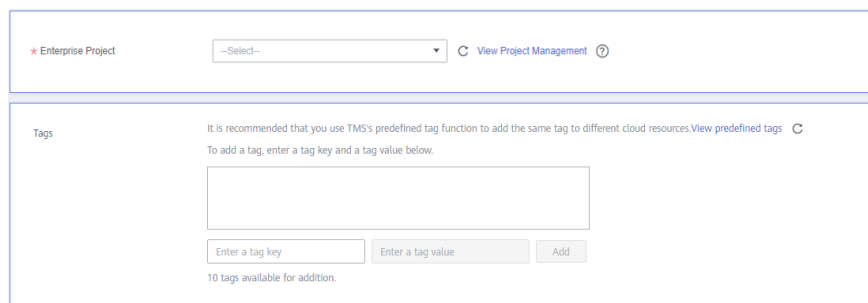


Table 4-6 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

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 not enabled, 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 not enabled, 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

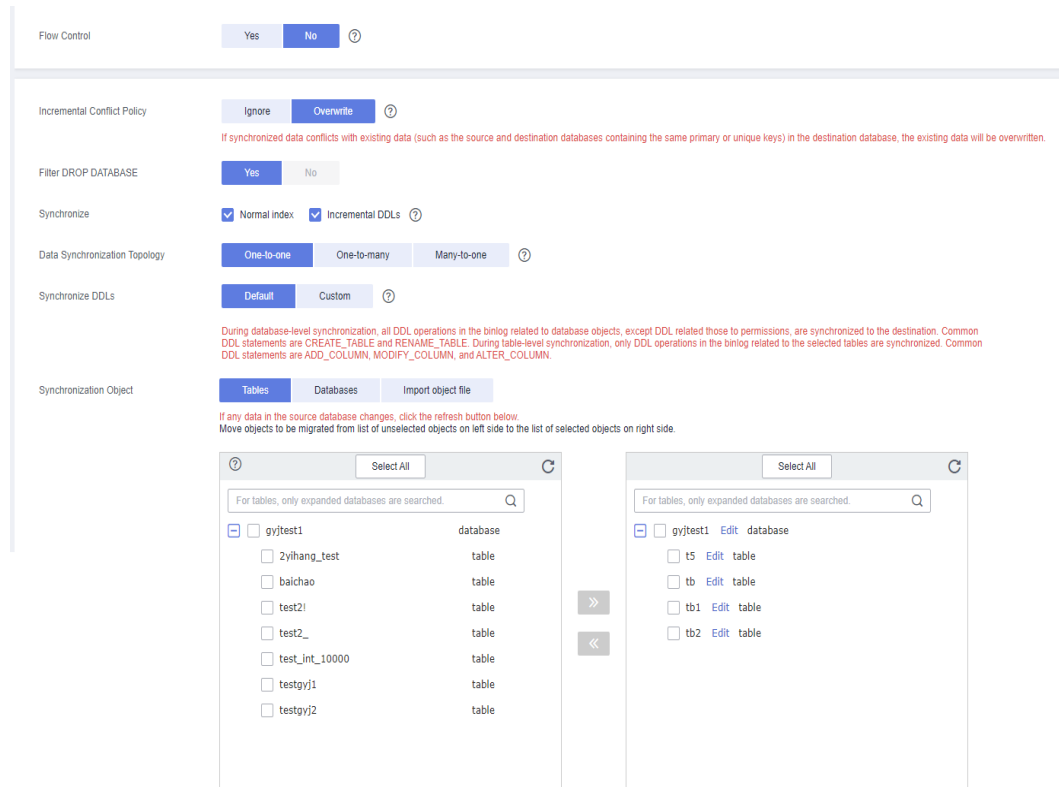
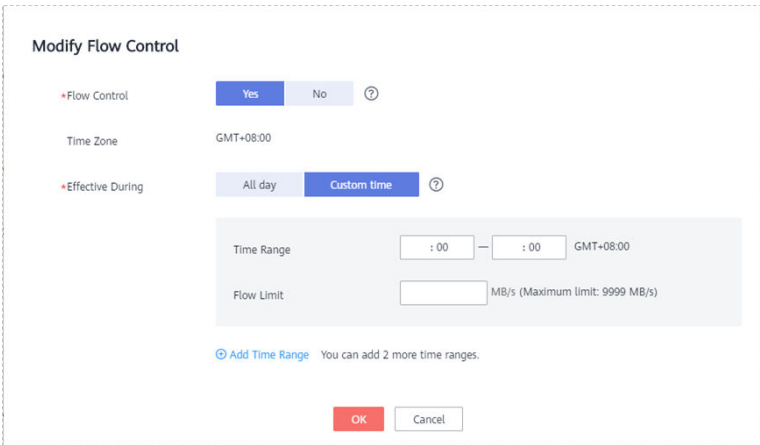



Table 4-9 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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. ● 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 data based on service requirements.</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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</p>

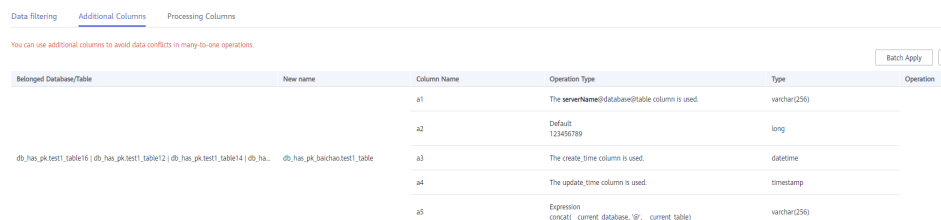
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: 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. <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 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 Column**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-9 Processing Data

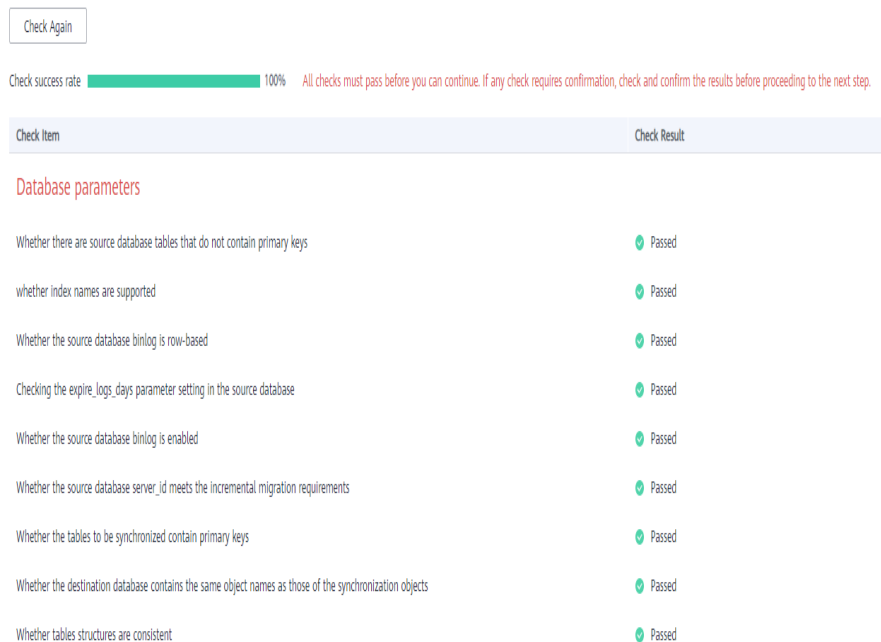


Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
		a1	The serverName@database@table column is used.	varchar(256)	
		a2	Default 123456789	long	
db_has_pk_test1_table16 db_has_pk_test1_table12 db_has_pk_test1_table14 db_has_pk_test1_table18		a3	The create_time column is used.	datetime	
		a4	The update_time column is used.	timestamp	
		a5	Expression concat(_current_database, '_')	varchar(256)	

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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-10 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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

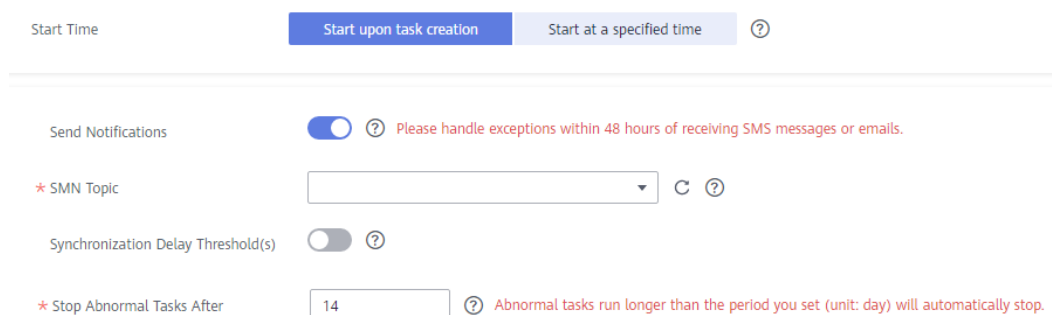



Table 4-10 Task startup settings

Parameter	Description
Started 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>SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.</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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.2 From MySQL to Kafka

Supported Source and Destination Databases

Table 4-11 Supported databases

Source DB	Destination DB
RDS for MySQL	Kafka

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](#).

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 [What Is the Impact of DRS on 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 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.

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, REPLICATION SLAVE, REPLICATION CLIENT, and RELOAD.
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 can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> • 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. • The database and table names in the source database cannot contain non-ASCII characters, or special characters '<'>/'\ • 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.

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. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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.) ● 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. ● During synchronization of table-level objects, renaming tables is not recommended.

Type	Restrictions
	<ul style="list-style-type: none"> • 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

Table 4-13 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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


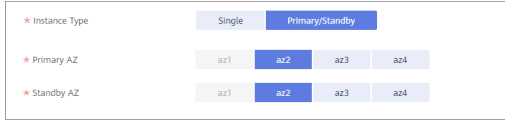


Table 4-14 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select Kafka .
Network Type	The Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
Instance Type	DRS instance type. The value can be Single or Primary/Standby . <ul style="list-style-type: none"> – Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. – Single: The single-node deployment architecture is used, which is cost-effective. This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover .
Source DB Instance	The 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 Subnet 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 ensure that there are available IP addresses. 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.
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-14 AZ</p> 

- DRS instance specifications

Figure 4-15 Specifications



Table 4-15 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-16 Enterprise projects and tags




Table 4-16 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.


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

✔ Test successful

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.
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 not enabled, 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 or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

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**.

Figure 4-19 Synchronization mode

Flow Control Yes No ?

Synchronize DML Insert Update Delete ?

Topic Synchronization Policy

Topic C

Synchronize Topic To

Data Format in Kafka JSON JSON-C

Synchronization Object Tables Databases Import object file

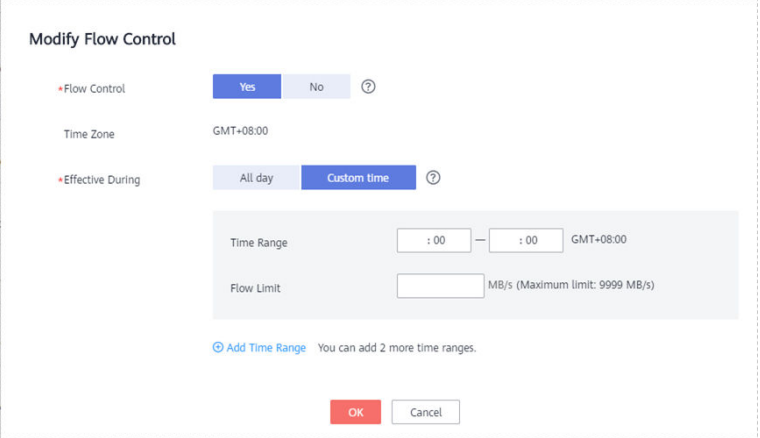
Only selected tables are synchronized. To synchronize new tables added to the source database during the incremental synchronization, you need to edit this synchronization task to select the new tables. If any data in the source database changes, click the refresh button below. Move objects to be migrated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

For tables, only expanded databases are searched.


- baichao database
- db_taurus_onsprecheck_001 database
- dengyuhuan database
- dengyuhuan1 database
- drs database
- drs4 database
- dyh1 database
- dyh1209 database

For tables, only expanded databases are searched.

Table 4-19 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</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 start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</p>

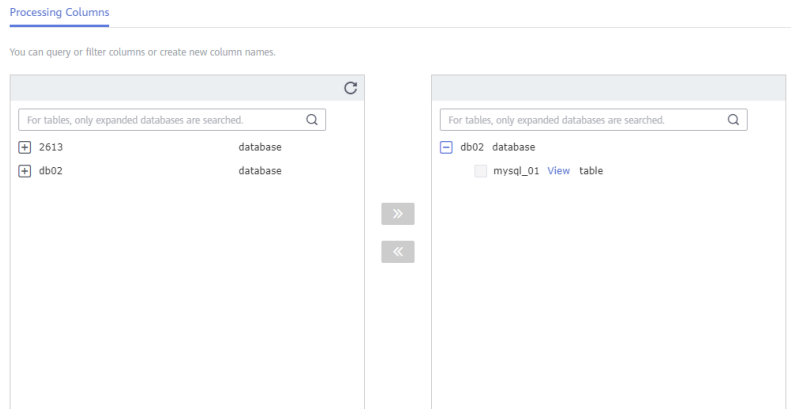
Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<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>
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.

Parameter	Description
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 when Synchronization Mode is set to Incremental in Step 2. ● 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>
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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. This function is available when **Synchronization Mode** is set to **Incremental** in [Step 2](#).

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 4-21 Processing data



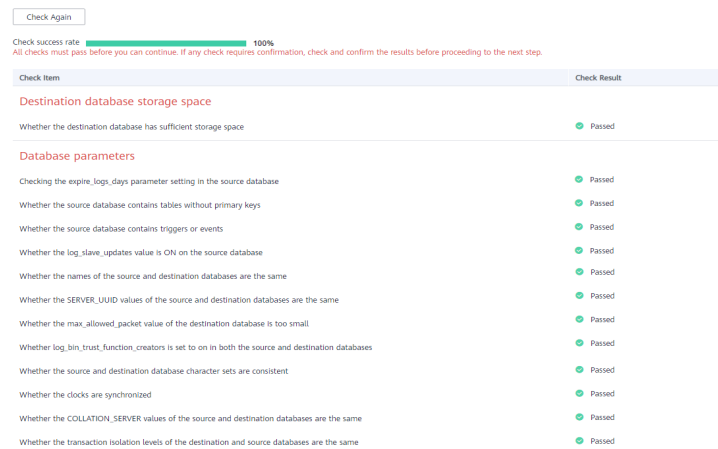
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-22 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----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

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

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
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.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● 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 ibttmp1. ● 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. ● 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. ● 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. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page." Below this, the form includes:

- Region:** A dropdown menu with a location pin icon.
- Project:** A dropdown menu.
- * Task Name:** A text input field containing "DRS-6131" and a help icon.
- Description:** A text area with a help icon and a character count "0/256" at the bottom right.

Table 4-23 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-25 Synchronization instance details

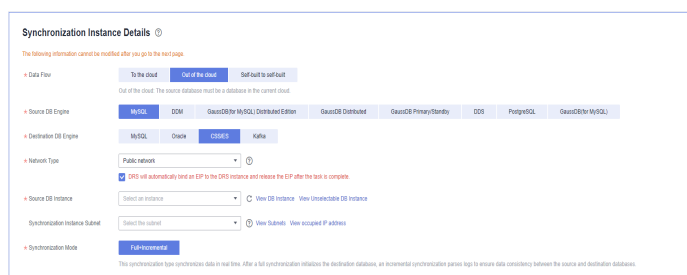


Table 4-24 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select CSS/ES .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source DB Instance	The RDS for MySQL instance you created.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-26 Specifications



Table 4-25 Specifications

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications .

- Enterprise Projects and Tags

Figure 4-27 Enterprise projects and tags

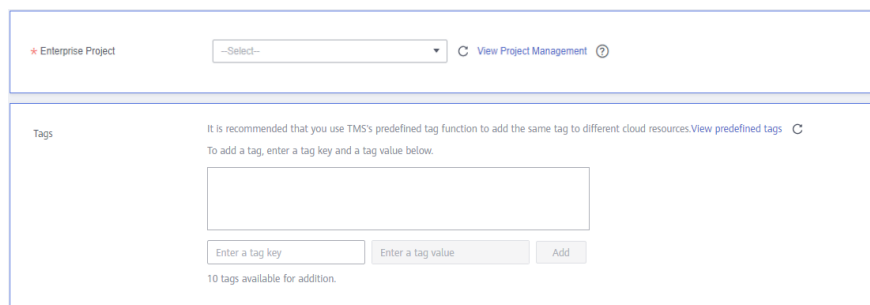


Table 4-26 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.


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

✔ Test successful

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 not enabled, 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

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

Encryption Certificate

Table 4-28 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.

Parameter	Description
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 not enabled, 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

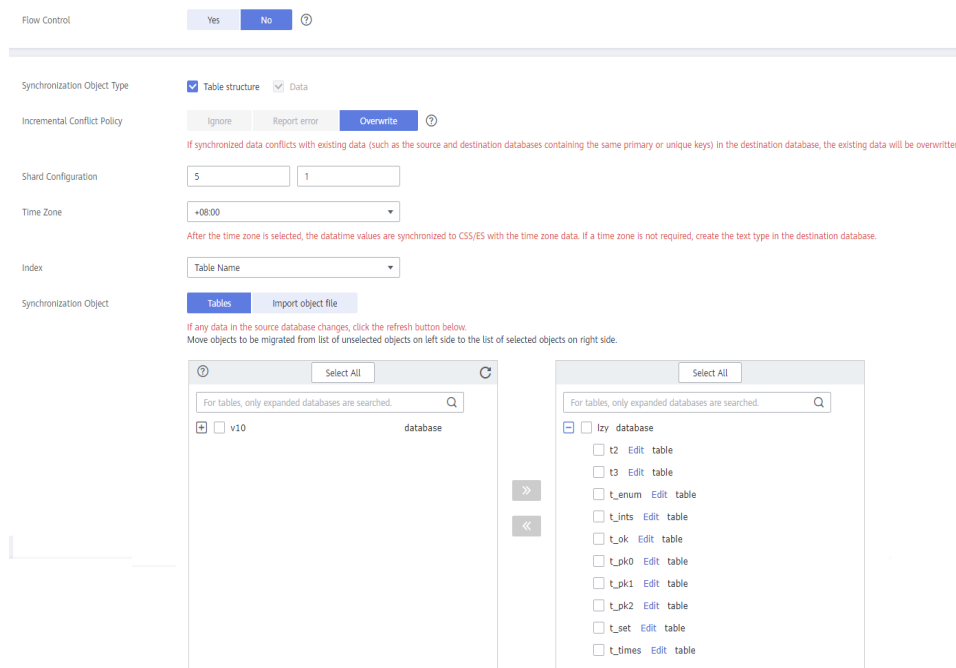
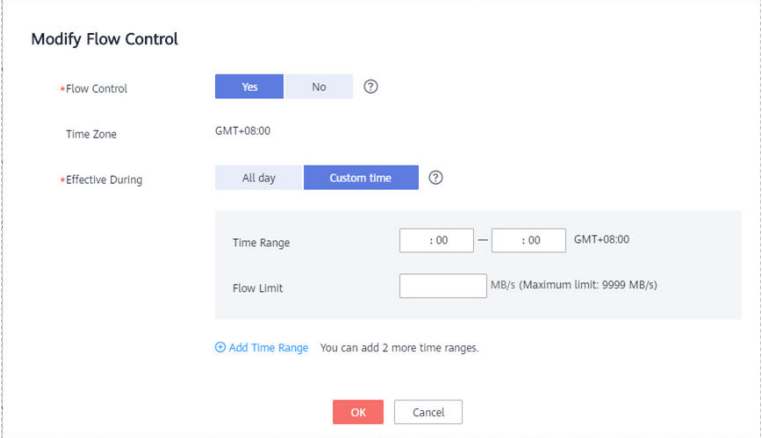



Table 4-29 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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.

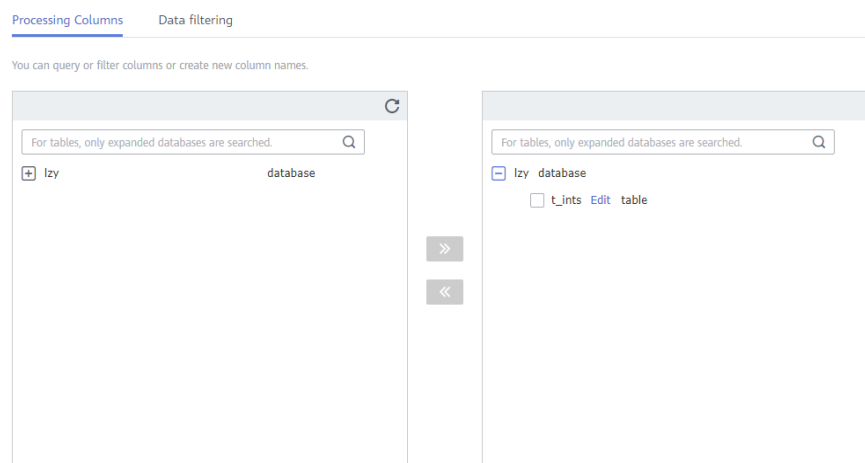
Parameter	Description
Synchronization Object Type	<p>Available options: Table structure and Data</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • 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>
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 the time zone is selected, the data of the datetime type will be synchronized to CSS/ES with the selected time zone. If the time zone is not required, you need to manually create the document type on the target end.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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



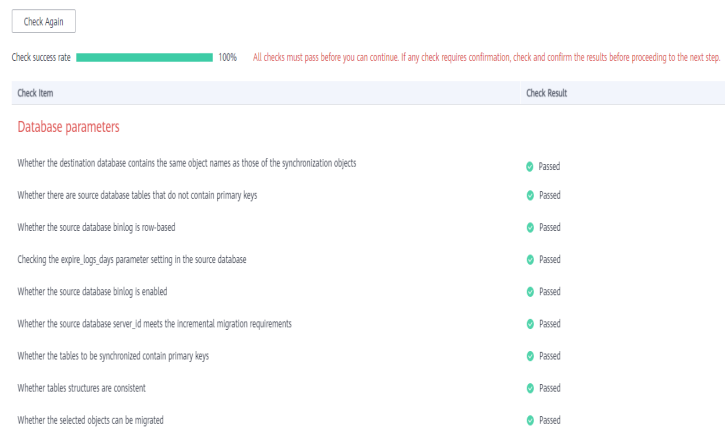
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-33 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-34 Task startup settings

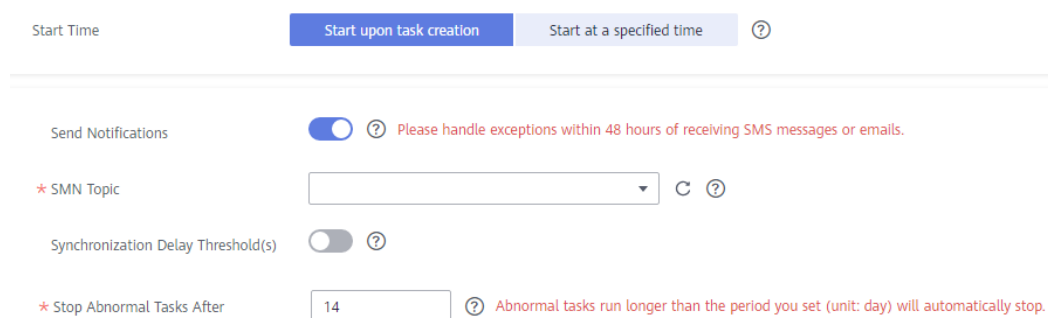



Table 4-30 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----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 	<ul style="list-style-type: none"> On-premises databases ECS databases

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 4-32 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.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● 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 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. ● 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. ● 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.
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.● 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.● 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 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.● 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.

Type	Restrictions
	<ul style="list-style-type: none"> • 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. • 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 4-35 Synchronization task information

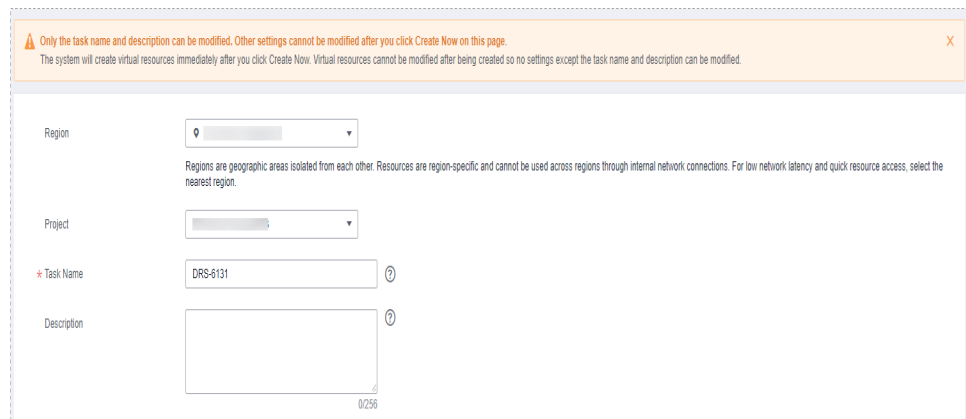


Table 4-33 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-36 Synchronization instance details

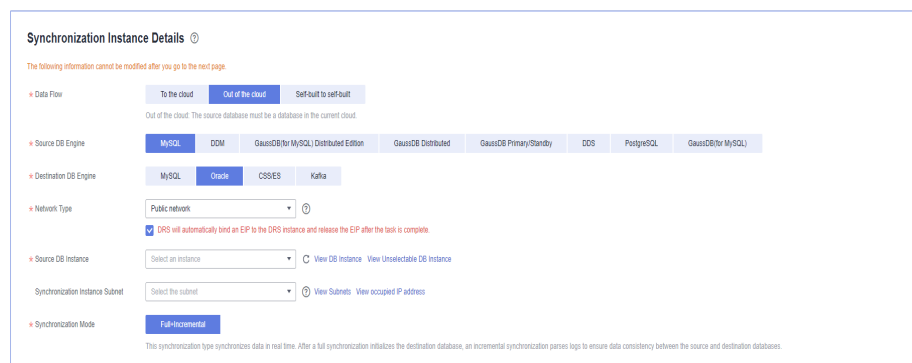


Table 4-34 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .

Parameter	Description
Source DB Engine	Select MySQL .
Destination DB Engine	Select Oracle .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source 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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-37 Specifications



Table 4-35 Specifications

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications .

- Enterprise Projects and Tags

Figure 4-38 Enterprise projects and tags

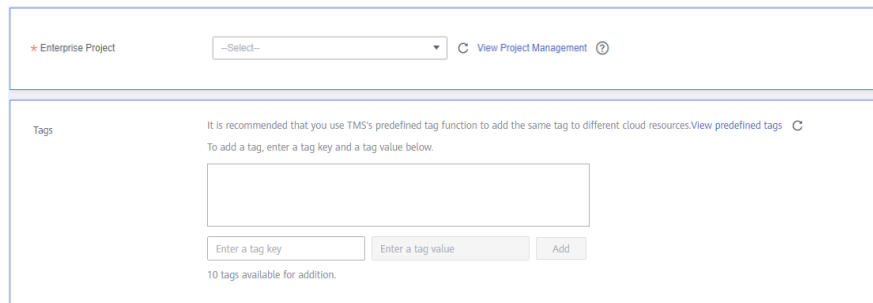


Table 4-36 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.


- 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 4-39 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection


 Test successful

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 not enabled, 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 4-40 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 4-38 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 not enabled, 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 4-41 Synchronization mode

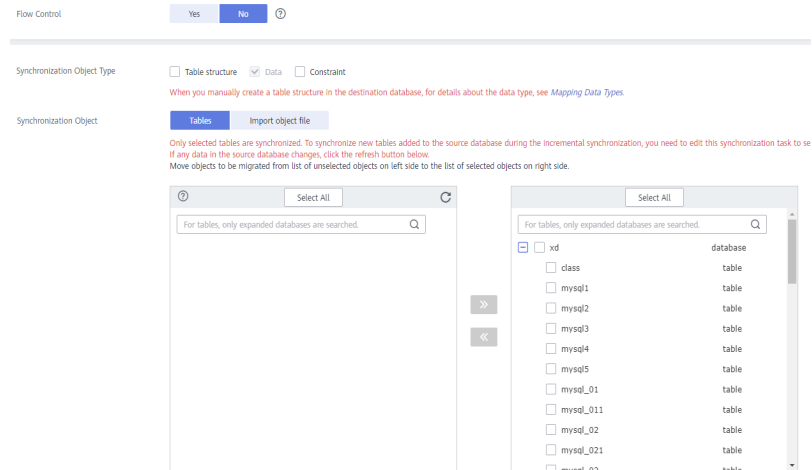
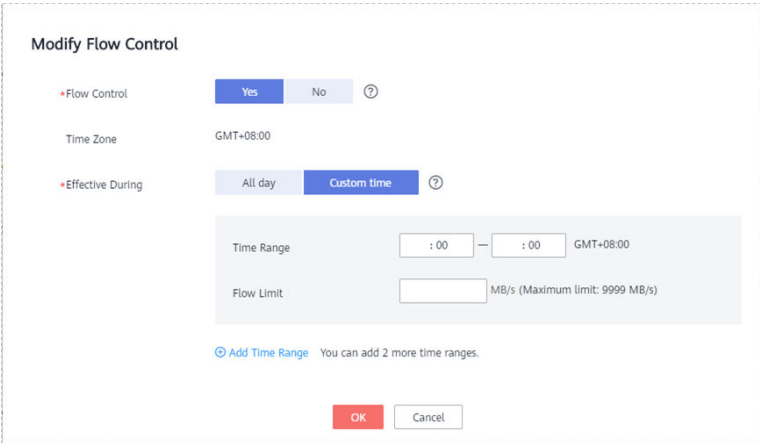



Table 4-39 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-42 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>Available options: Table structure, Data, and Constraint</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.
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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

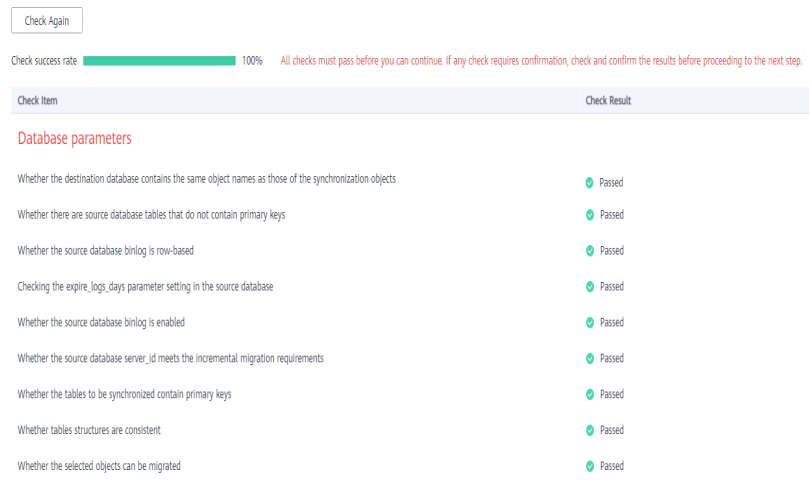
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-43 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-44 Task startup settings

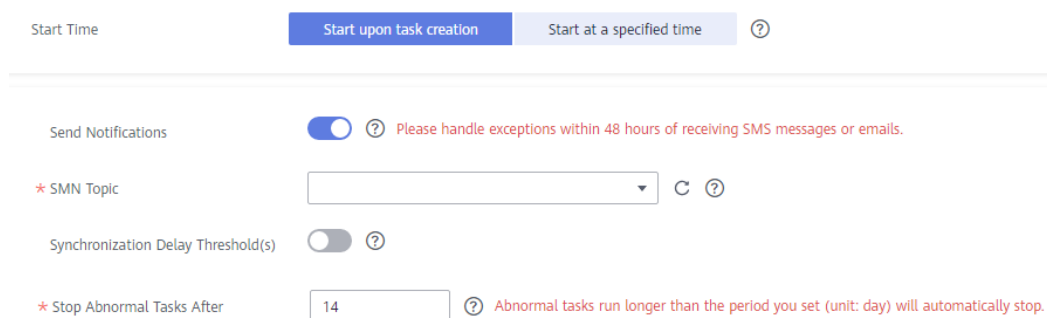



Table 4-40 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.5 From DDM to MySQL

Supported Source and Destination Databases

Table 4-41 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">DDM instances	<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECSMySQL databases on other cloudsRDS for MySQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 4-42 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 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: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, and UPDATE. The root account of the RDS for MySQL DB instance has the preceding permissions by default.
Synchronization object	<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 timestamp. • 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.

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 source database binlog 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 the following characters: '<>\' and non-ASCII characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● Enable GTID of the source database.
Destination database	<ul style="list-style-type: none"> ● The destination database is an on-premises MySQL database. ● The destination DB instance must have sufficient storage 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. ● 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.

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 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 destination DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● 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. ● 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. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key when the server system breaks down. ● 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. ● 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. ● During an incremental synchronization, do not perform the restoration operation on the source database.

Type	Restrictions
	<ul style="list-style-type: none"> ● During incremental synchronization, 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 '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. - 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. ● During incremental synchronization, perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute.

4.6 From DDM to Oracle

Supported Source and Destination Databases

Table 4-43 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

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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

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 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.

Table 4-44 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 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.

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 database sharding middleware cannot contain the following characters: '<>\' and non-ASCII characters. ● 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. ● 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 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 unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. ● 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.

Type	Restrictions
	<ul style="list-style-type: none">• During an incremental synchronization, do not perform the restoration operation on the source database.

4.7 From DDM to Kafka

Supported Source and Destination Databases

Table 4-45 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• DDM instances	<ul style="list-style-type: none">• Kafka

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](#).

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 [What Is the Impact of DRS on 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 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.

Table 4-46 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 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.
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 the following characters: '<>\' and non-ASCII characters. 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.

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. • 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.

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 4-45 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. X
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: [Text input: DRS-0131] ⓘ

Description: [Text area] ⓘ

0/255

Table 4-47 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-46 Synchronization instance details

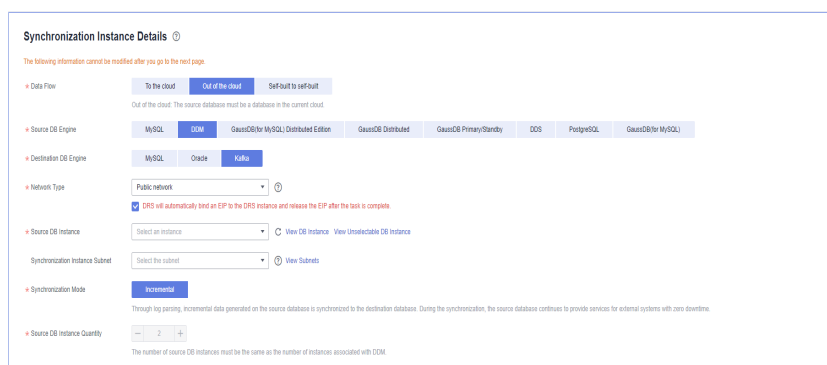


Table 4-48 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDM .
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
Source DB Instance	An available DDM instance.

Parameter	Description
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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 DB Instance Quantity	Specifies the number of DB instances bound to the source DDM database. Set this parameter based on the site requirements.

- DRS instance specifications

Figure 4-47 Specifications



Table 4-49 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-48 Enterprise projects and tags

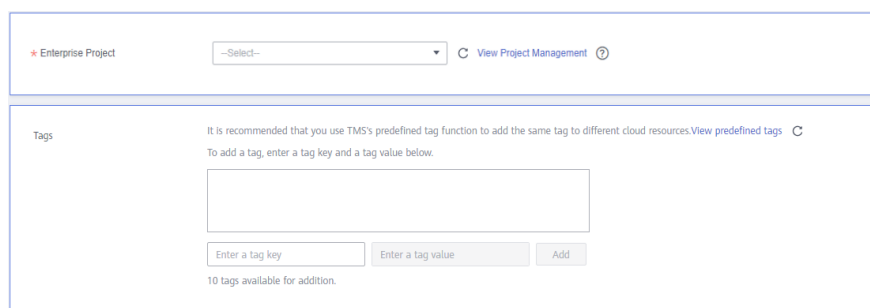


Table 4-50 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-49 Source database information

Source Database

DB Instance Name rds-00w-source

Database Username

Database Password

✔ Test successful

Table 4-51 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 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 4-50 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.


Security Protocol 

Table 4-52 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 .

 **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 4-51 Synchronization mode

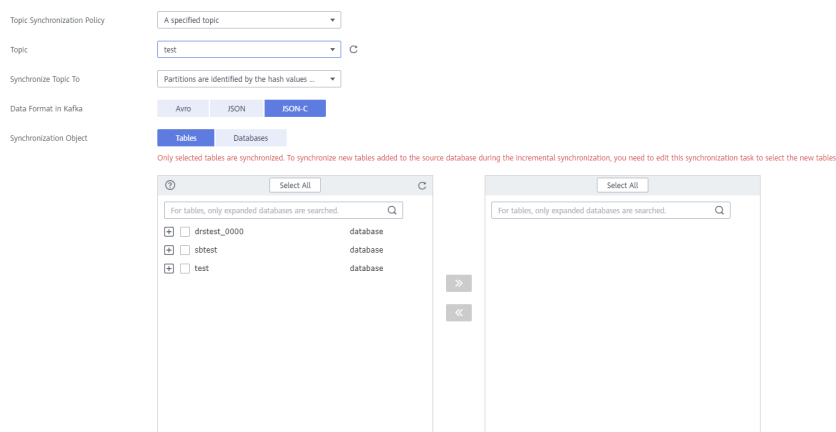



Table 4-53 Synchronization Object

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 .

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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

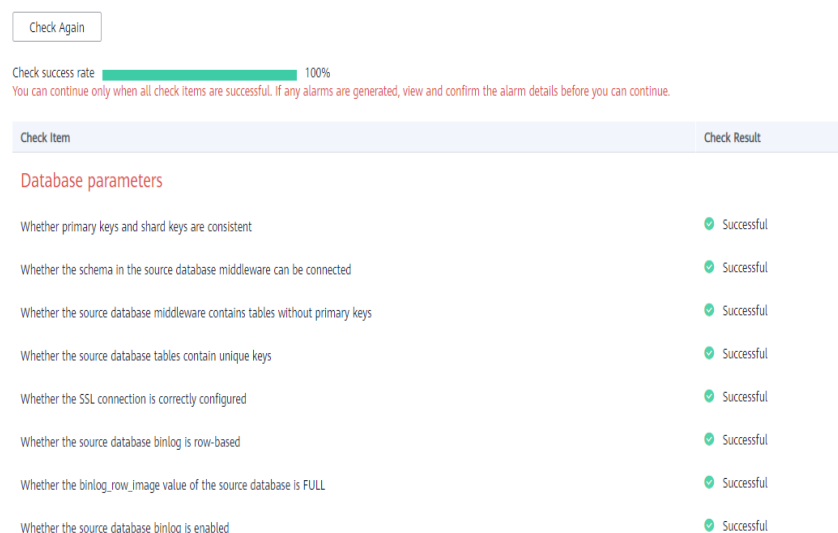
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-52 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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

The screenshot shows the 'Task startup settings' configuration page. At the top, there are two tabs for 'Start Time': 'Start upon task creation' (selected) and 'Start at a specified time'. Below this, there are four settings:


- Send Notifications:** A toggle switch is turned on. A red note says: "Please handle exceptions within 48 hours of receiving SMS messages or emails."
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the number '14'. A red note says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-54 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.8 From DDS to MongoDB

Supported Source and Destination Databases

Table 4-55 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DDS instances 	<ul style="list-style-type: none"> • On-premises MongoDB databases • MongoDB databases on an ECS • MongoDB database on other clouds

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](#).

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 [What Is the Impact of DRS on 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 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.

Table 4-56 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>
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 <code>2</code>. ● 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.
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 <code>\/. \$</code> or spaces. The collection name and view name cannot start with <code>system.</code> or contain the dollar sign (<code>\$</code>).

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 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 4-54 Synchronization task information

The screenshot shows a form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "+ Task Name" (a text input field containing "DRS-6131"), and "Description" (a text area with a character count of "0/256").

Table 4-57 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-55 Synchronization instance details

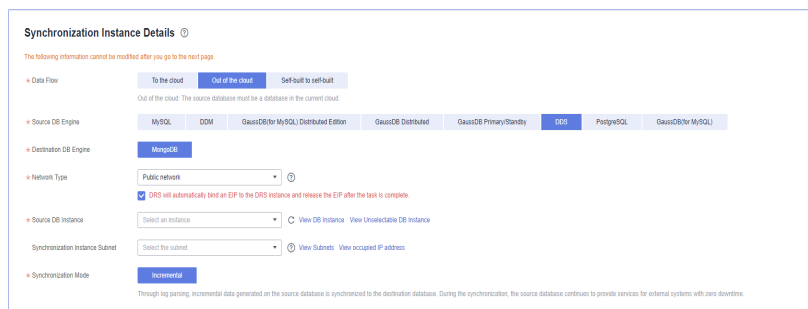


Table 4-58 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select DDS.
Destination DB Engine	Select MongoDB .
Network Type	The Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
Source DB Instance	An available DDS instance
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-56 Specifications



Table 4-59 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-57 Enterprise projects and tags

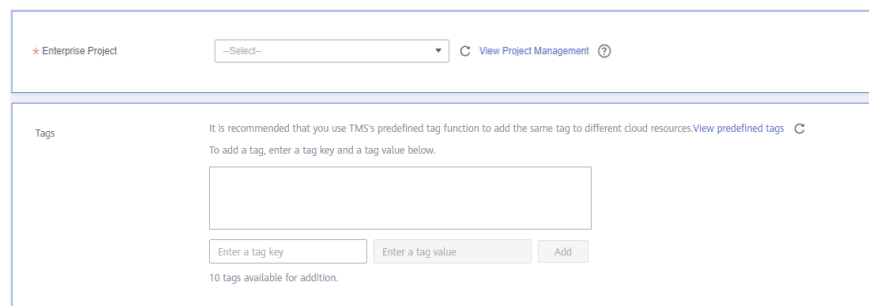


Table 4-60 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> – This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. – 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 ends.

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 4-58 Source database information

Source Database

DB Instance Name

Database Username


Database Password 

Table 4-61 Source database settings


Parameter	Description
DB Instance Name	The DDS instance that you selected when creating the replication instance and 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 4-59 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

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

Encryption Certificate

Table 4-62 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 not enabled, 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 4-60 Synchronization mode

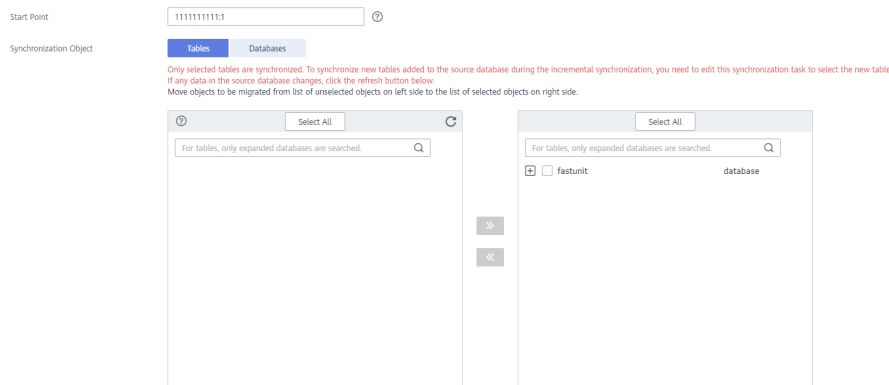



Table 4-63 Synchronization Object

Parameter	Description
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.
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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

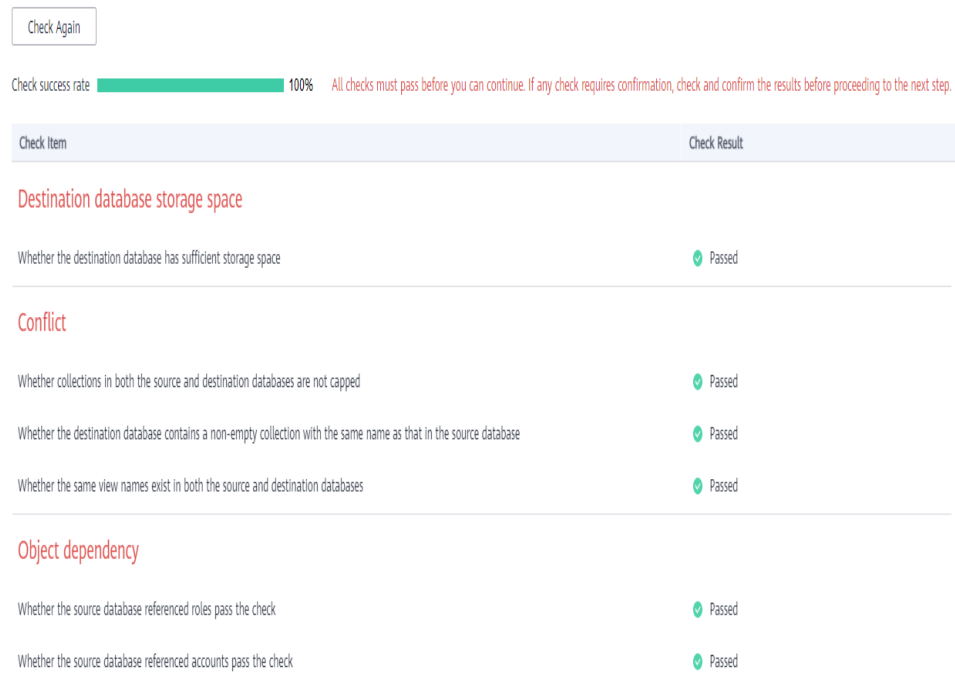
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-61 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-62 Task startup settings

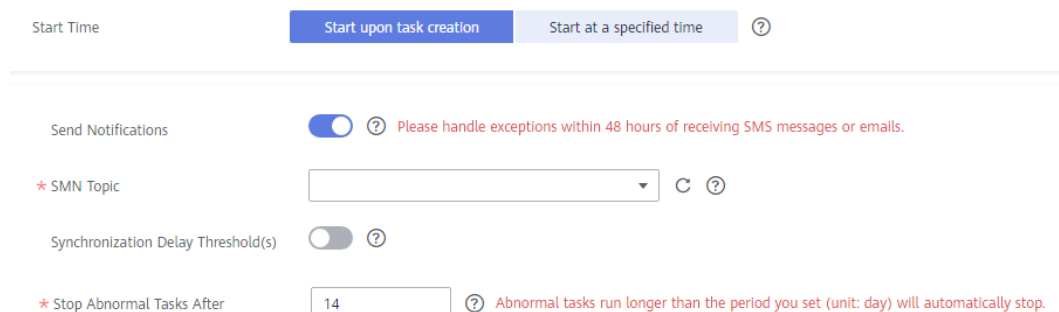



Table 4-64 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.9 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 4-65 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 4-66 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-66 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. 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>CAUTION 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>

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 4-67](#). 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.

Table 4-67 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 [What Is the Impact of DRS on 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 4-68 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. ● 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. - 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. • 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.
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.

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 4-63 Synchronization task information

Table 4-69 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-64 Synchronization instance details

Table 4-70 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
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
Source DB Instance	Specifies an RDS PostgreSQL DB instance you have created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-65 Specifications



Table 4-71 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-66 Enterprise projects and tags




Table 4-72 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 4-67 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-73 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 4-68 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Connection Method

✔ Test successful

Table 4-74 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 4-69 Synchronization mode

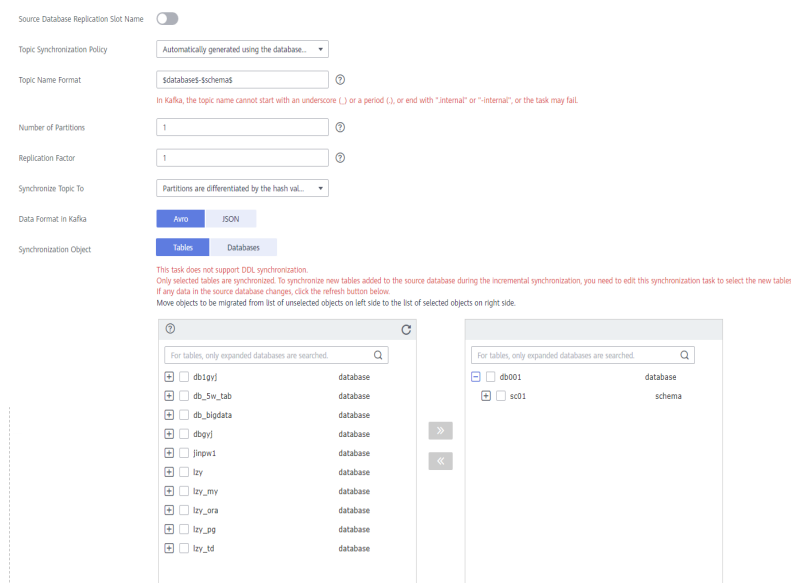



Table 4-75 Synchronization object

Parameter	Description
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 .

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, 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- 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

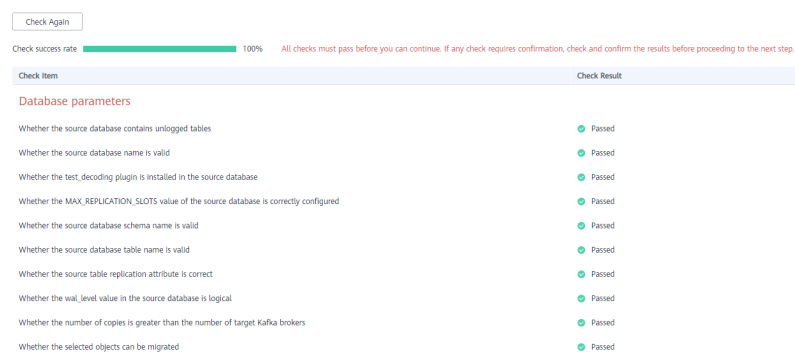
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-70 Task 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-71 Task startup settings

The screenshot shows the 'Task startup settings' configuration interface. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below the tabs, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says: "Please handle exceptions within 48 hours of receiving SMS messages or emails."
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the number '14'. A red note says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-76 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.10 From GaussDB Distributed to MySQL

Supported Source and Destination Databases

Table 4-77 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 4-78](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-78 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: INTEGER, TINYINT, SMALLINT, BIGINT, 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 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase.- 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 4-79](#). 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 4-79 Database user permission

Type	Incremental
Source database user	The user must have the sysadmin role or the following minimum permissions: <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.
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.
 - Start your synchronization task during off-peak hours. 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.
 - 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.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 4-80 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <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 or GB18030 character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data synchronization may be inconsistent, 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 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">- 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.- During real-time synchronization, the consistency of distributed transactions is not ensured.- The table structure information is saved in uppercase in the source database. During synchronization, if the table names in the destination database are different from those in the source database, map the source table names to the destination table names.- 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	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. • 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 <i>GaussDB Developer Guide</i>.
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.
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.

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 4-72 Synchronization task information

Table 4-81 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-73 Synchronization instance details

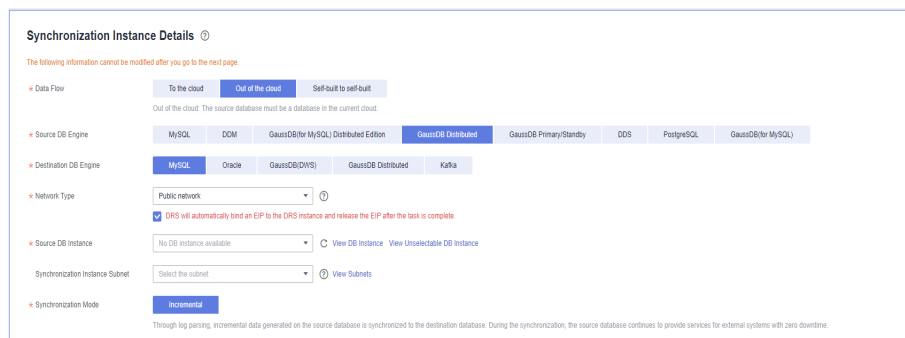


Table 4-82 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	– Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

- Enterprise Projects and Tags

Figure 4-74 Enterprise projects and tags




Table 4-83 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-75 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-84 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 4-76 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 4-85 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 not enabled, 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 4-77 Synchronization mode

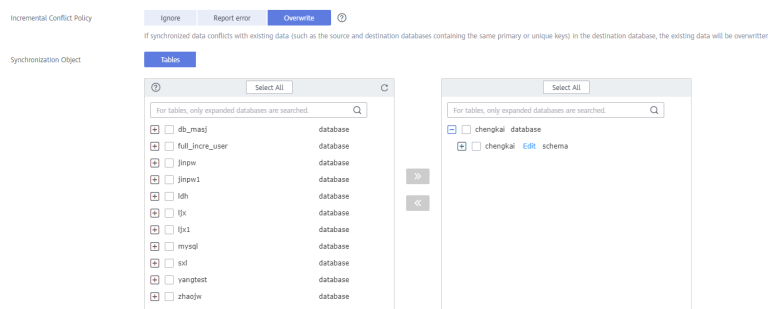



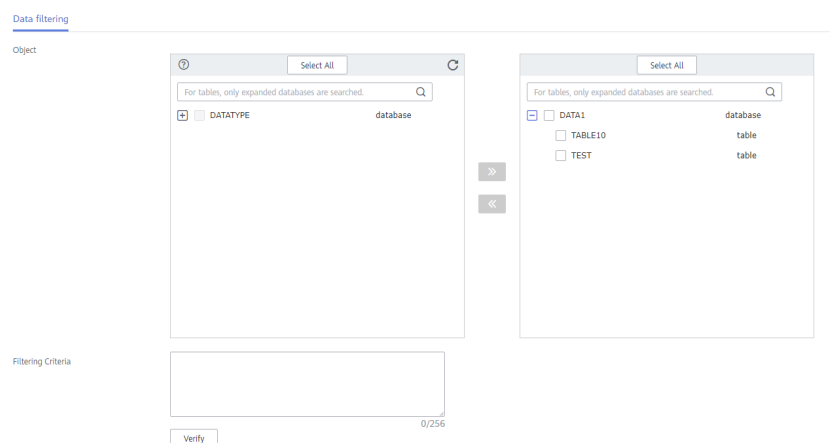
Table 4-86 Synchronization Object

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. ● 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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](#).

Figure 4-78 Processing data



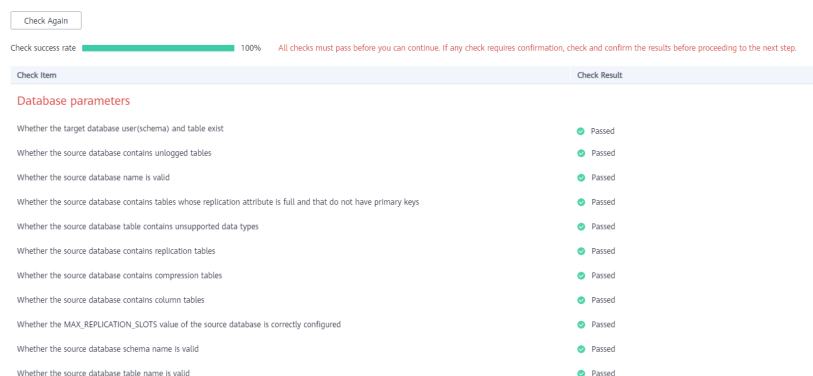
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-79 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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

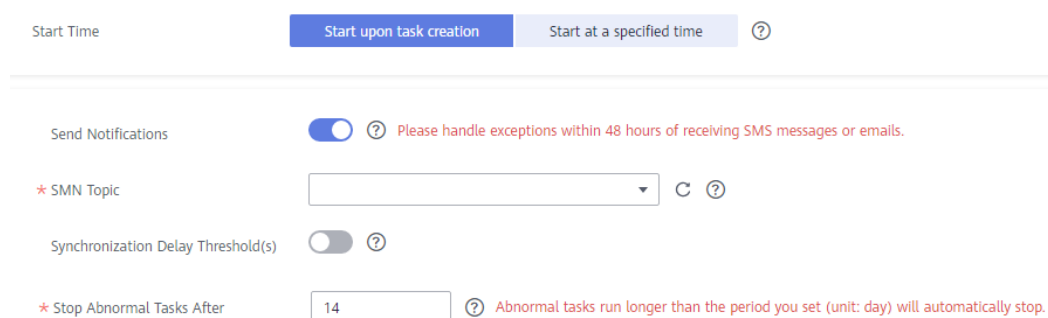



Table 4-87 Task startup settings

Parameter	Description
Started 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>SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.</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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.11 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 4-88 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 4-89 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-89 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-90](#). 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 4-90 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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.

Type	Full	Incremental	Full+Incremental
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

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-91 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 or GB18030 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 synchronization may be inconsistent, 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"> - 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. - During real-time synchronization, the consistency of distributed transactions is not ensured. - 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. - 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.
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.

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. • 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 <i>GaussDB Developer Guide</i>. • 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 an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task 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.

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 4-81 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "* Task Name" (a text input field containing "DRS-6131"), and "Description" (a text area with a character count of "0/256").

Table 4-92 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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

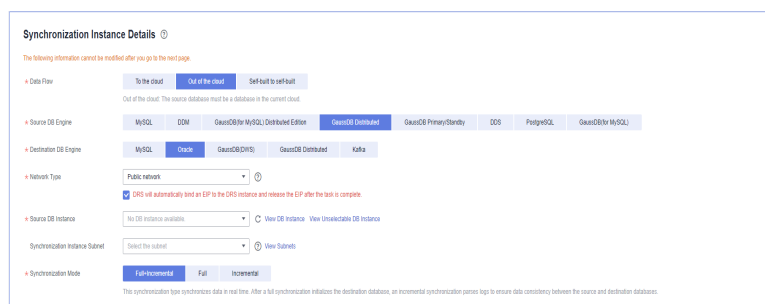


Table 4-93 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Oracle .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 4-83 Enterprise projects and tags

Table 4-94 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-84 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-95 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 4-85 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 4-96 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 not enabled, 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 4-86 Synchronization mode


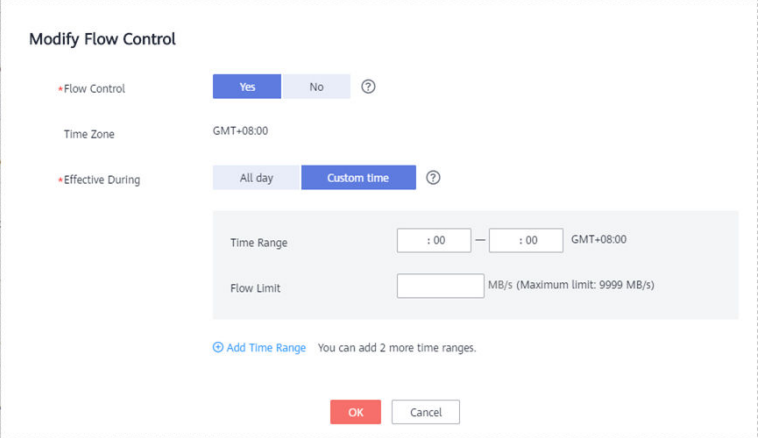



Table 4-97 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-87 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. ● 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 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 4-88 Parameter settings

Full Synchronization Settings

* Concurrent Export Tasks ?

* Concurrent Import Tasks ?

* Rows per Shard ?

Incremental Replay Settings

* Concurrent Replay Tasks ?

Table 4-98 Full synchronization settings

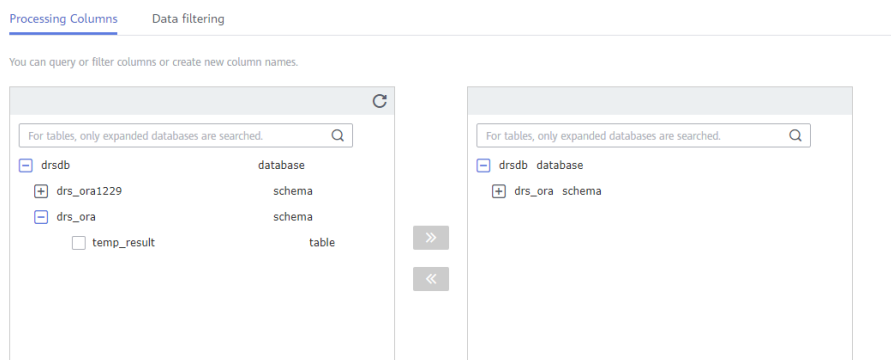
Parameter	Description	Default Value
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

Table 4-99 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

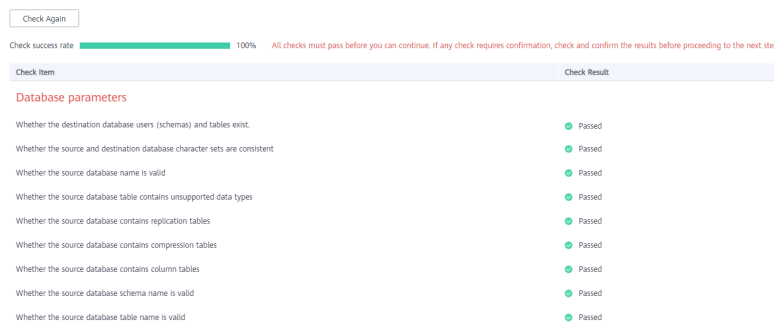
- Step 6** On the **Process Data** page, set the filtering rules for data processing.
- If data processing is not required, click **Next**.
 - If you need process columns or filter out data, set the corresponding rules by referring to **Processing Data**.

Figure 4-89 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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-90 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 8** On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-91 Task startup settings





Table 4-100 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.12 From GaussDB Distributed to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-101 Supported databases

Source DB	Destination DB
GaussDB distributed	GaussDB(DWS) cluster

Supported Synchronization Objects

Table 4-102 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-102 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-103](#). 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 4-103 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	The user has the sysadmin role or the following permissions: CONNECT permission for DATABASE, USAGE permission for SCHEMA, and SELECT or UPDATE (required only for tables without primary keys) permission for tables.	The user has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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.
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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-104 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 on the current cloud. - 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.
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.

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 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 <i>GaussDB Developer Guide</i>.
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	<ul style="list-style-type: none"> • Stop a task normally. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed.

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 4-92 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-6131 [Text input]

Description: [Text area]

0/256

Table 4-105 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-93 Synchronization instance details

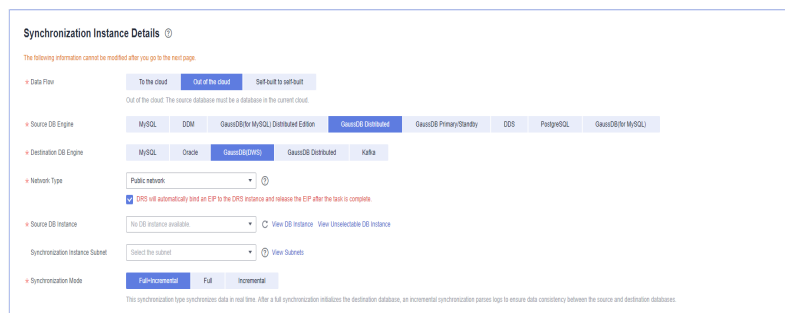


Table 4-106 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

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.

- DRS instance specifications

Figure 4-94 Specifications



Table 4-107 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-95 Enterprise projects and tags

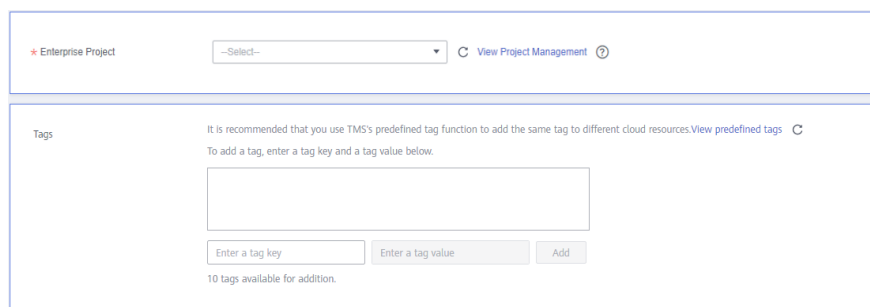


Table 4-108 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-96 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-109 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 4-97 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

Table 4-110 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 4-98 Synchronization mode

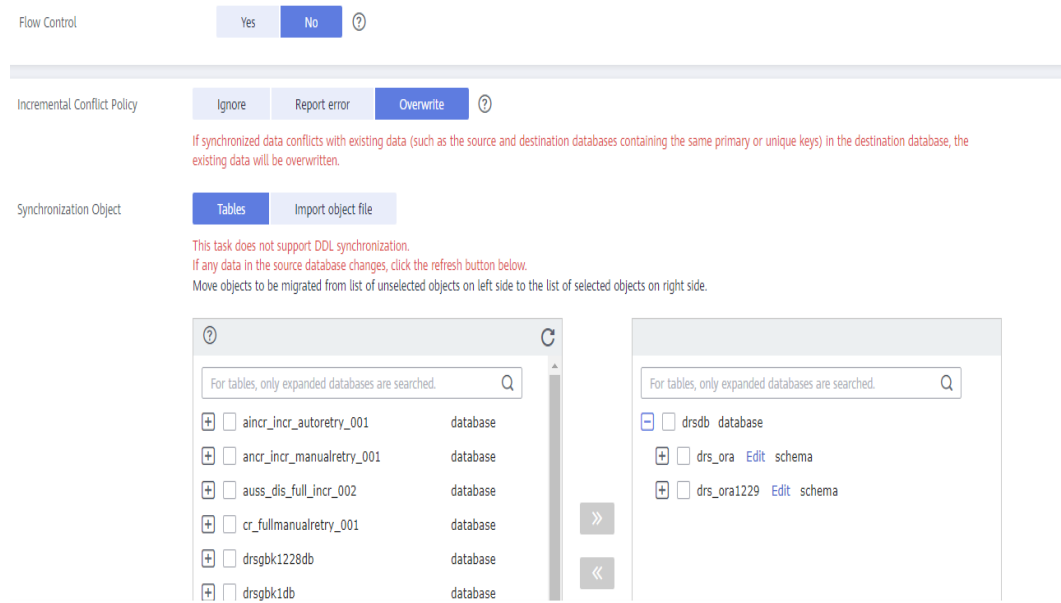
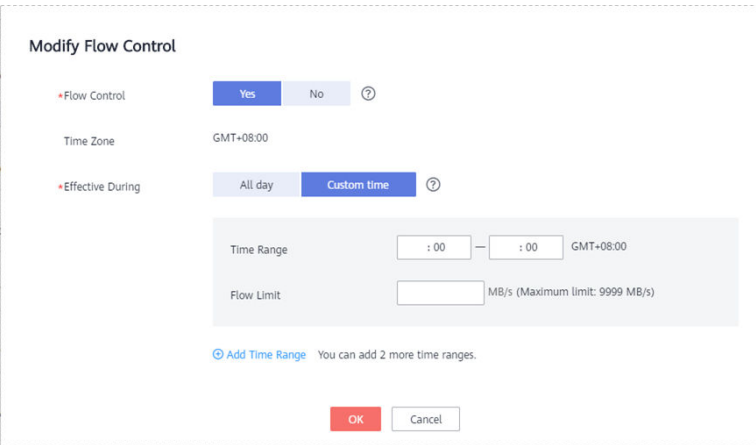



Table 4-111 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-99 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.


Figure 4-100 Parameter settings

Full Synchronization Settings

* Synchronization Object Type Table structure Data Constraint(excluding foreign keys)

* Concurrent Export Tasks 8 

* Concurrent Import Tasks 8 

* Rows per Shard 

Incremental Replay Settings


* Concurrent Replay Tasks 64 

Table 4-112 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

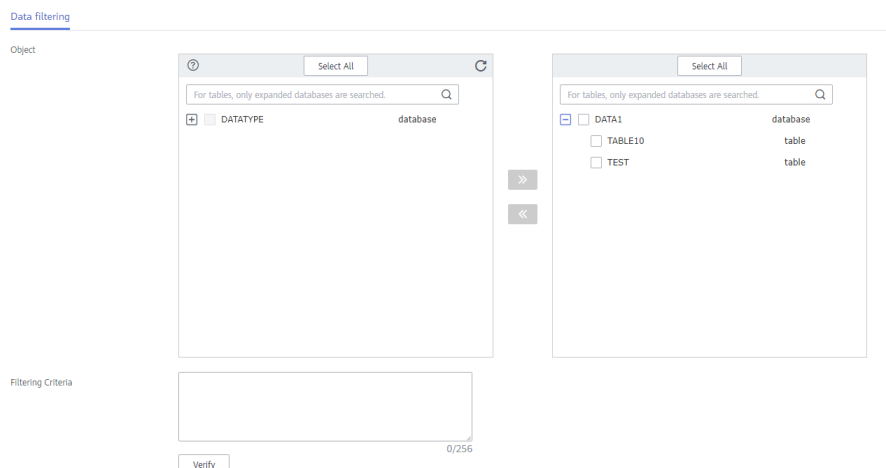
Table 4-113 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

Figure 4-101 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-102 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-103 Task startup settings

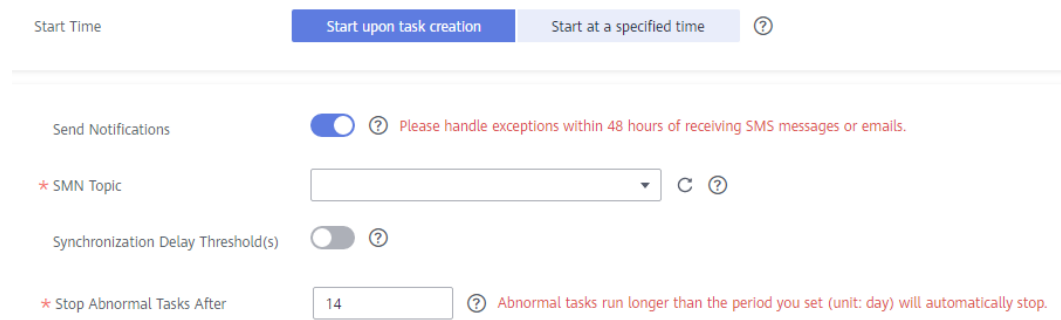



Table 4-114 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.13 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 4-115 Supported databases

Source DB	Destination DB
GaussDB distributed	Kafka 0.11 or later

Supported Synchronization Objects

[Table 4-116](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-116 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 scenario: Incremental synchronization • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, 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"> - 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - 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 4-117](#). 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 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 4-117 Database user permission

Type	Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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.

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 [What Is the Impact of DRS on 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 4-118 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. ● Other notes: <ul style="list-style-type: none"> - The source must be a GaussDB distributed instance on Huawei Cloud. - 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.

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. • 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 <i>GaussDB Developer Guide</i>.
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> – When an 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 an incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed.

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 4-104 Synchronization task information

Table 4-119 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-105 Synchronization instance details

Table 4-120 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .

Parameter	Description
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Kafka .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-106 Specifications



Table 4-121 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-107 Enterprise projects and tags

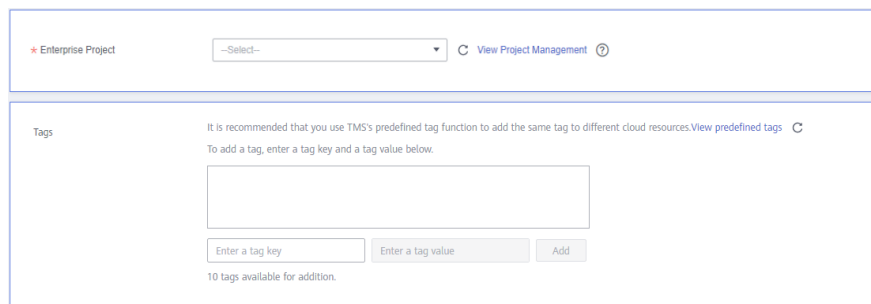


Table 4-122 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 4-108 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-123 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 4-109 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Connection Method

✔ Test successful

Table 4-124 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 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 4-110 Synchronization mode





Table 4-125 Synchronization Object

Parameter	Description
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 .

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, 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-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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces.

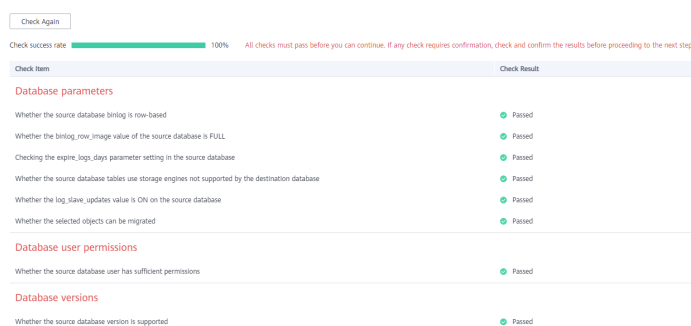
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-111 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-112 Task startup settings

The screenshot shows the 'Task startup settings' configuration interface. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below the tabs, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says: "Please handle exceptions within 48 hours of receiving SMS messages or emails."
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the number '14'. A red note says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 4-126 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.14 From GaussDB Distributed to GaussDB Distributed

Supported Source and Destination Databases

Table 4-127 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 4-128 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-128 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-129](#). 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 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 4-129 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. 	<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. 	<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.

Type	Full	Incremental	Full+Incremental
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, and INDEX permissions for tables. 	<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, and INDEX permissions for tables are required. To compare content, the MONADMIN permission is required. 	<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, and INDEX permissions for tables are required. To compare content, the MONADMIN permission is 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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-130 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"> - The destination must be a GaussDB distributed instance on Huawei Cloud. - 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. - 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. - During real-time synchronization, the consistency of distributed transactions is not ensured. - 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.
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.

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. • 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 <i>GaussDB Developer Guide</i>.
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	<ul style="list-style-type: none"> • Stop a task normally. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed.

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 4-113 Synchronization task information

Table 4-131 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-114 Synchronization instance details

Table 4-132 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .

Parameter	Description
Source DB Engine	Select GaussDB Distributed .
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
Source DB Instance	The GaussDB distributed instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full+Incremental , Full , and Incremental . Full+Incremental is used as an example. - 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.

- Enterprise Projects and Tags

Figure 4-115 Enterprise projects and tags




Table 4-133 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-116 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-134 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 4-117 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

Table 4-135 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 4-118 Synchronization mode

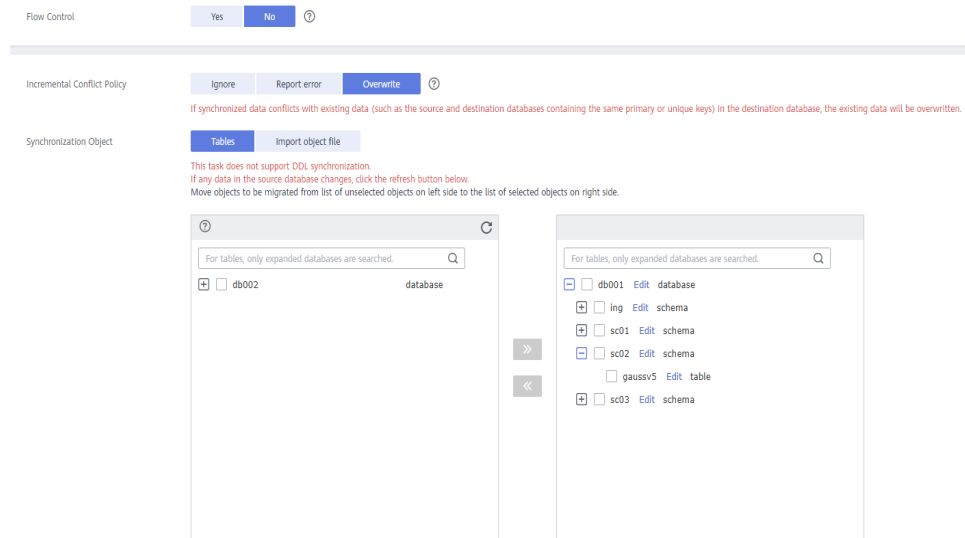
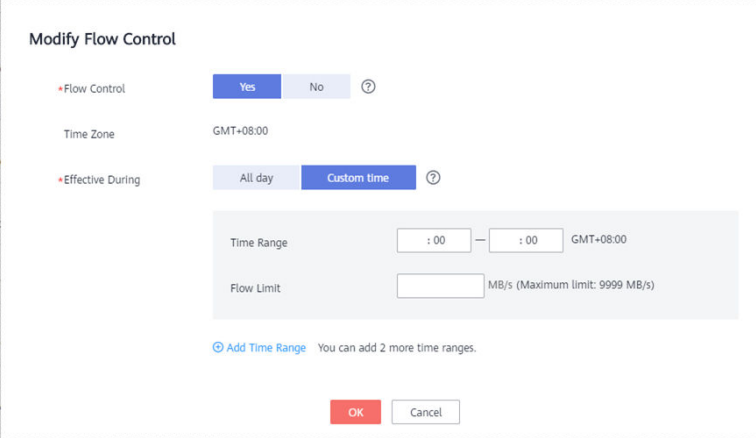



Table 4-136 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-119 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>Ignore and overwrite: Synchronization stability is prioritized, so tasks will not be interrupted as data conflicts occur.</p> <p>Report error: Data quality is prioritized. Any data conflicts are not allowed, so once a conflict occurs, the synchronization task fails and an error is reported. You need to manually find the cause of the fault. If the task is in the failed state for a long time, the storage space may be used up and the task cannot be restored.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in **Step 2** and click **Next**.

Figure 4-120 Parameter settings

Full Synchronization Settings

* Synchronization Object Type Table structure Data Constraint

* Concurrent Export Tasks 8 (?)

* Concurrent Import Tasks 8 (?)

* Rows per Shard (?)

Incremental Replay Settings

* Concurrent Replay Tasks 64 (?)

Table 4-137 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8

Parameter	Description	Default Value
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> • 0 All tables are not sharded, and each table is synchronized as a whole. • Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

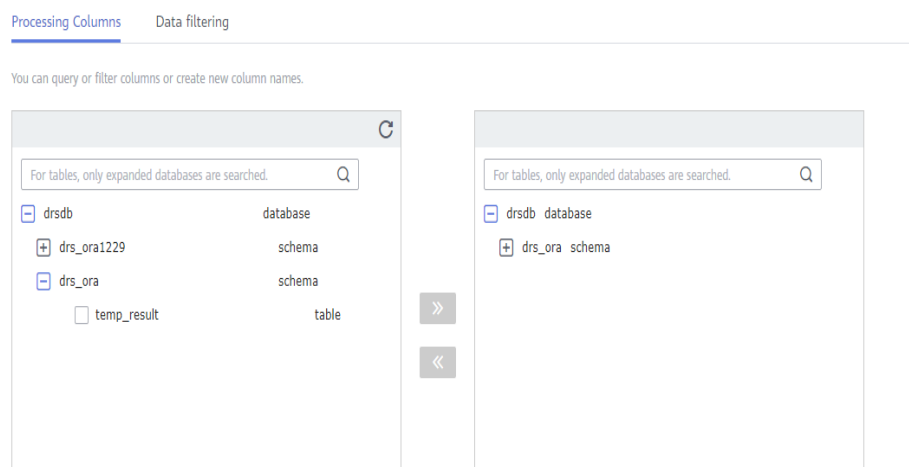
Table 4-138 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

Figure 4-121 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-122 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 8** On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-123 Task startup settings

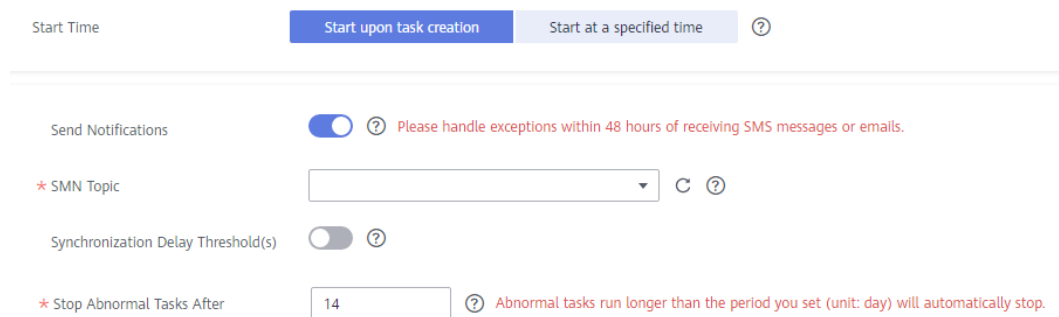



Table 4-139 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

Parameter	Description
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>
Synchronization Delay Threshold	<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 Notification.
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> <p>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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.15 From GaussDB Distributed to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-140 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 4-141](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-141 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-142](#). 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 4-142 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. 	<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. 	<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.

Type	Full	Incremental	Full+Incremental
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, and INDEX permissions for tables. 	<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, and INDEX permissions for tables. 	<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, and INDEX permissions for 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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-143 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"> - The destination must be a GaussDB primary/standby instance on Huawei Cloud. - 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. - 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. - During real-time synchronization, the consistency of distributed transactions is not ensured. - 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.
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.

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 <i>GaussDB Developer Guide</i>.
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.
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. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted. Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed.

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 4-124 Synchronization task information

Table 4-144 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-125 Synchronization instance details

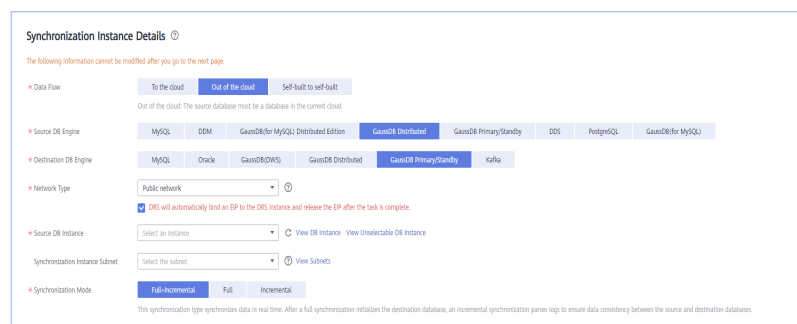


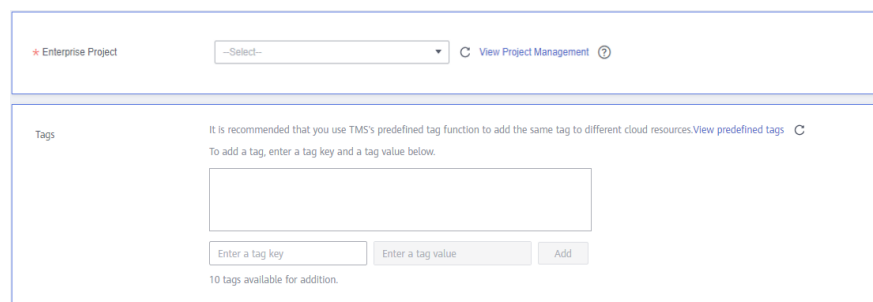
Table 4-145 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select GaussDB Primary/Standby .
Network Type	The public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 4-126 Enterprise projects and tags



The screenshot displays a user interface for configuring enterprise projects and tags. At the top, there is a section for 'Enterprise Project' with a dropdown menu currently showing '--Select--' and a 'View Project Management' link. Below this is a 'Tags' section. It includes a large text area for notes, followed by two input fields: 'Enter a tag key' and 'Enter a tag value', and an 'Add' button. A note at the bottom of the tags section states '10 tags available for addition.'

Table 4-146 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-147 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 ?

Ensure that the entered addresses belong to the same DB instance

Database Username

Database Password

✔ Test successful

Table 4-148 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.
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 4-129 Synchronization mode

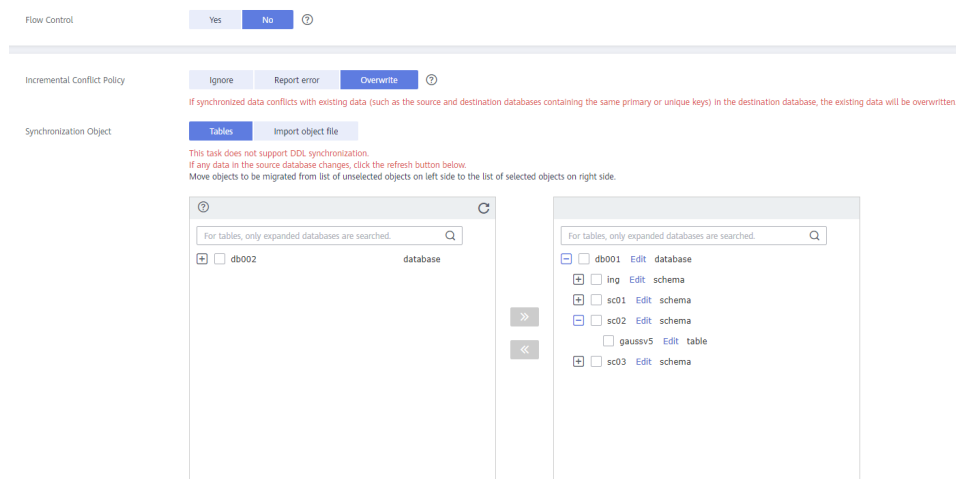
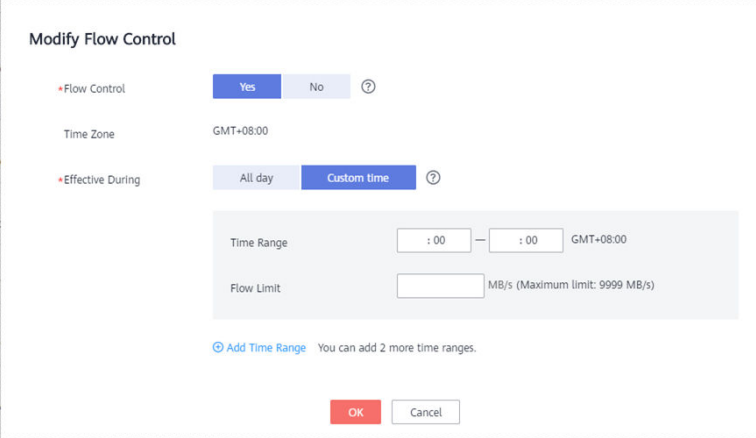



Table 4-149 Synchronization object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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.</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>Ignore and overwrite: Synchronization stability is prioritized, so tasks will not be interrupted as data conflicts occur.</p> <p>Report error: Data quality is prioritized. Any data conflicts are not allowed, so once a conflict occurs, the synchronization task fails and an error is reported. You need to manually find the cause of the fault. If the task is in the failed state for a long time, the storage space may be used up and the task cannot be restored.</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 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 + <code>_key</code> to prevent index/constraint name conflicts. The prefix of the index is <code>i_</code>, the prefix of the constraint is <code>c_</code>, 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in **Step 2** and click **Next**.

Figure 4-131 Parameter settings

Full Synchronization Settings

★ Synchronization Object Type Table structure Data Constraint

★ Concurrent Export Tasks

★ Concurrent Import Tasks

★ Rows per Shard

Incremental Replay Settings

★ Concurrent Replay Tasks

Table 4-150 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

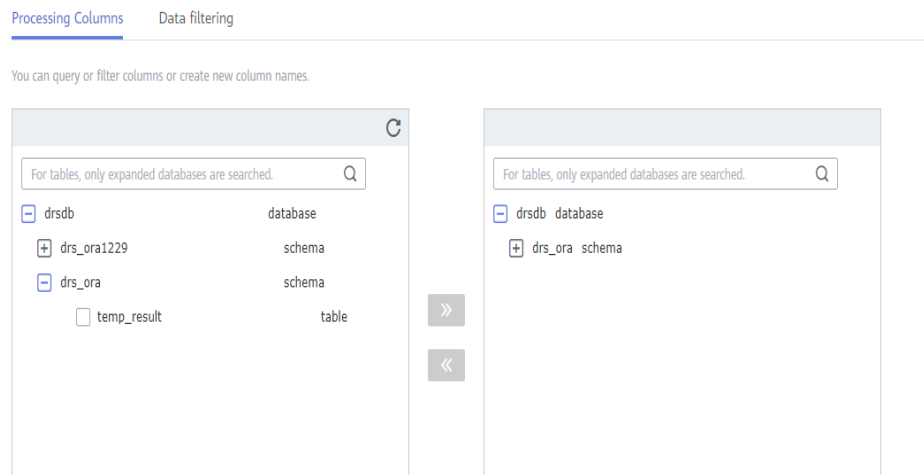
Table 4-151 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

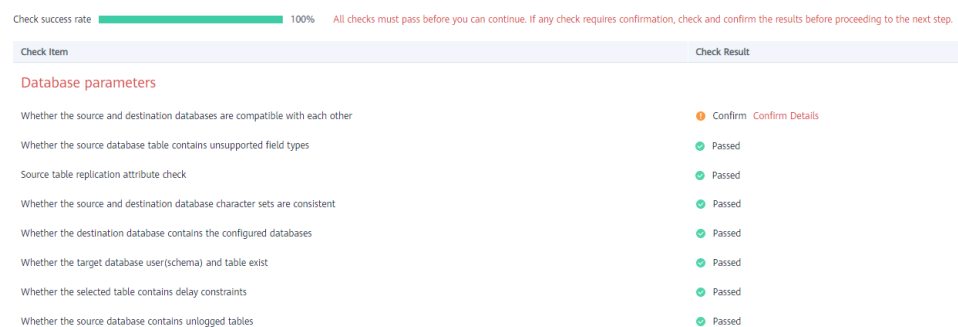
Figure 4-132 Processing data



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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-133 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-134 Task startup settings

The screenshot shows the 'Task startup settings' configuration page. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below the tabs, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says 'Please handle exceptions within 48 hours of receiving SMS messages or emails.'
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the value '14'. A red note says 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 4-152 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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
Synchronization Delay Threshold	<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 Notification.
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> <p>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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.16 From GaussDB Primary/Standby to MySQL

Supported Source and Destination Databases

Table 4-153 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

 **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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-155](#). 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.

Table 4-155 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>
Destination database user	<p>Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES</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.

- Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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	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 or GB18030 character set and the destination database uses the GBK, GB2312, or GB18030 character set, some characters may be incompatible due to database differences, some data may contain garbled characters, data synchronization may be inconsistent, 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. - 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. - The table structure information is saved in uppercase in the source database. During synchronization, if the table names in the destination database are different from those in the source database, map the source table names to the destination table names. - 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.
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.

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. • 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 <i>GaussDB Developer Guide</i>.
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.
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.

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 4-135 Synchronization task information

Table 4-157 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-136 Synchronization instance details

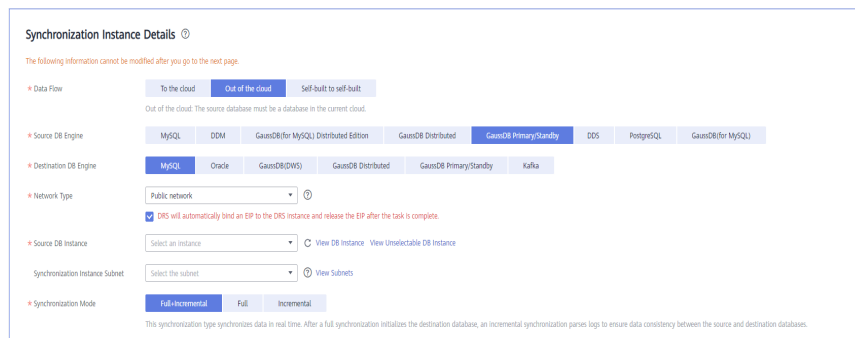


Table 4-158 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

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.

- Enterprise Projects and Tags

Figure 4-137 Enterprise projects and tags

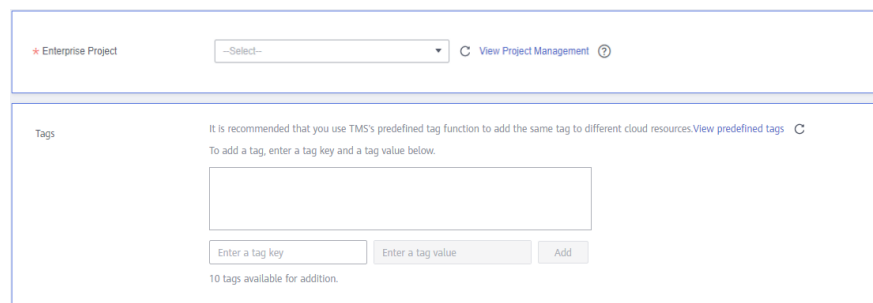


Table 4-159 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-138 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-160 Source 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 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 4-139 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection


 Test successful

Table 4-161 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.

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 not enabled, 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 4-140 Synchronization mode

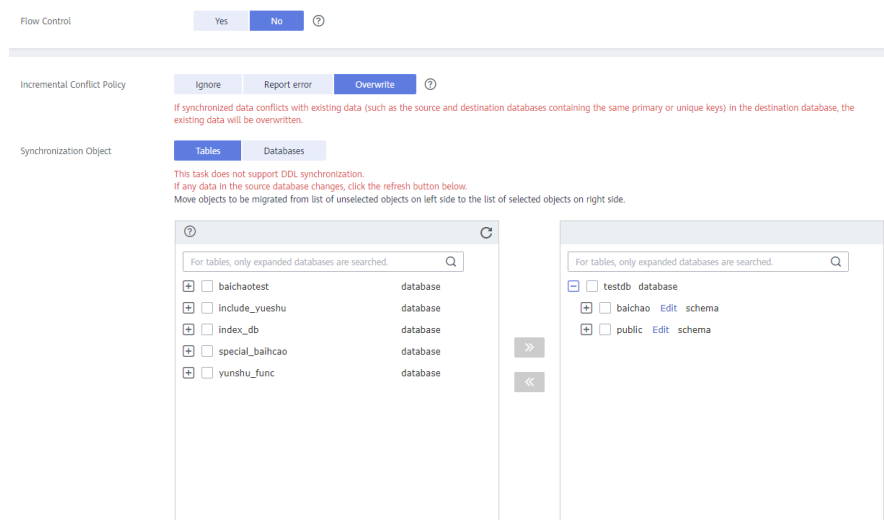
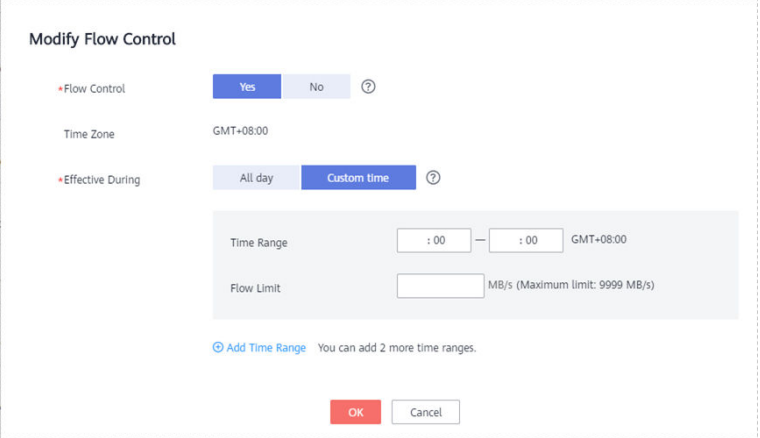



Table 4-162 Synchronization Object

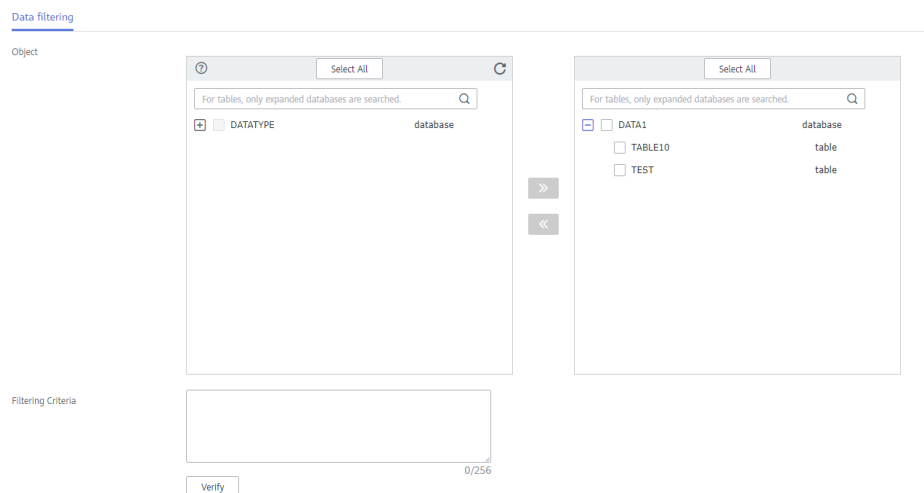
Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of 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 All day. A maximum of three 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 4-141 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. 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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](#).

Figure 4-142 Processing data



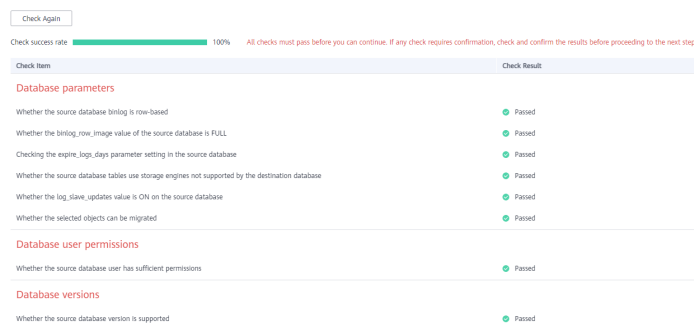
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-143 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-144 Task startup settings

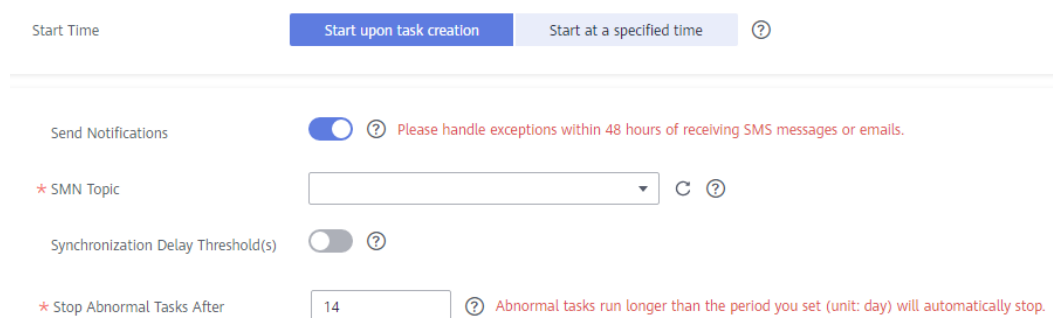



Table 4-163 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.17 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 4-164 Supported databases

Source DB	Destination DB
GaussDB primary/standby	<ul style="list-style-type: none">On-premises Oracle 11g and 19c databasesOracle 11g and 19c databases on an ECS

Supported Synchronization Objects

Table 4-165 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-165 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase.- 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 4-166](#). 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.

Table 4-166 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>
Destination database user	<p>The user must have the DBA role or the following minimum permissions:</p> <p>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:</p> <p>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:</p> <p>CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-167 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 or GB18030 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 synchronization may be inconsistent, 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"> - 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. - 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. - 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.
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.

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. ● 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 <i>GaussDB Developer Guide</i>. ● 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 and unique key 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 an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task 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.

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 4-145 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. X

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 ⓘ

Description: [Text area] ⓘ

0/256

Table 4-168 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-146 Synchronization instance details

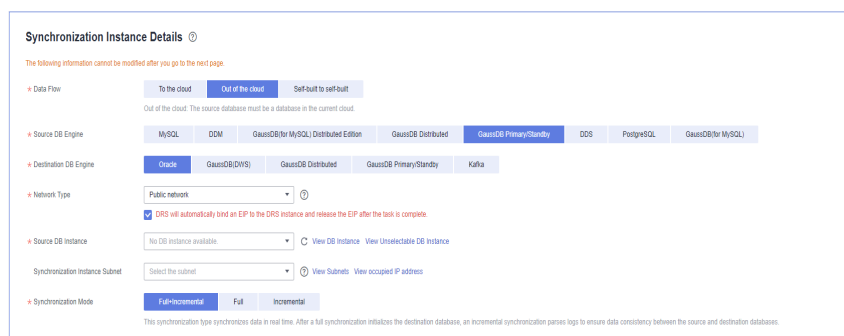


Table 4-169 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
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
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 4-147 Enterprise projects and tags

The screenshot displays a user interface for managing enterprise projects and tags. At the top, there is a section for 'Enterprise Project' with a dropdown menu currently showing '--Select--' and a 'View Project Management' link. Below this is a 'Tags' section. It includes a large text area for notes, followed by two input fields: 'Enter a tag key' and 'Enter a tag value', and an 'Add' button. A note at the bottom of the tags section states '10 tags available for addition'.

Table 4-170 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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">- This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- 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 ends.

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 4-148 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-171 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 4-149 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 4-172 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 not enabled, 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 4-150 Synchronization mode

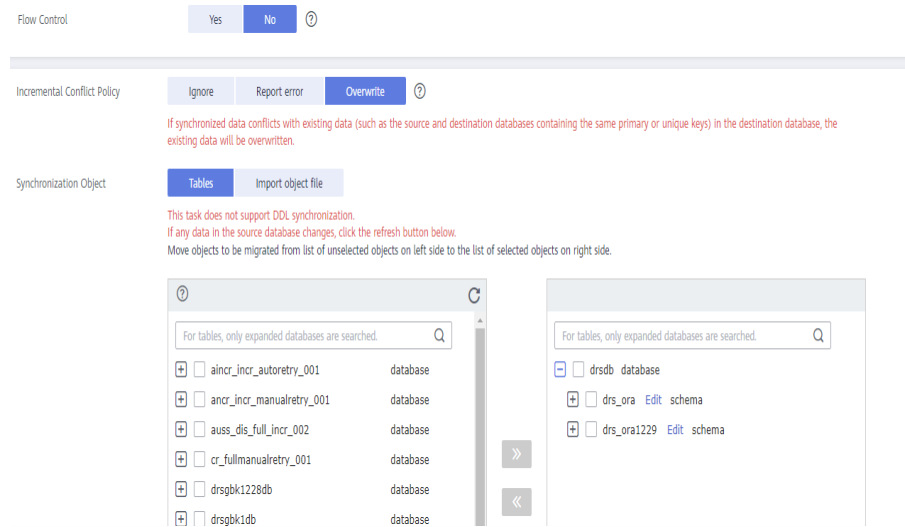
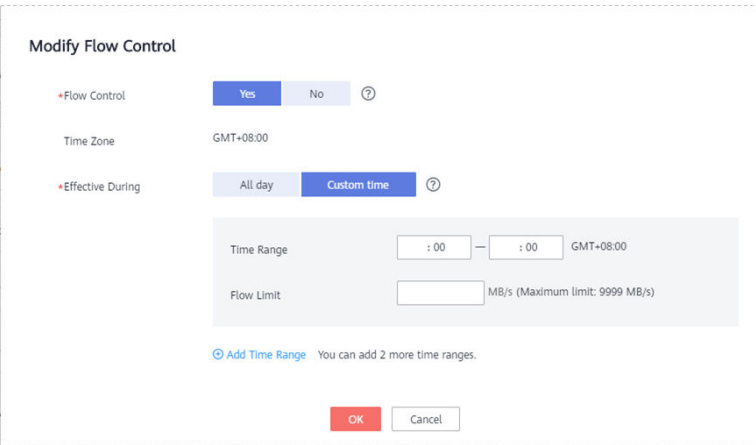



Table 4-173 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-151 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 4-152 Parameter settings

Full Synchronization Settings

* Concurrent Export Tasks ?

* Concurrent Import Tasks ?

* Rows per Shard ?

Incremental Replay Settings

* Concurrent Replay Tasks ?

Table 4-174 Full synchronization settings

Parameter	Description	Default Value
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

Table 4-175 Incremental synchronization settings

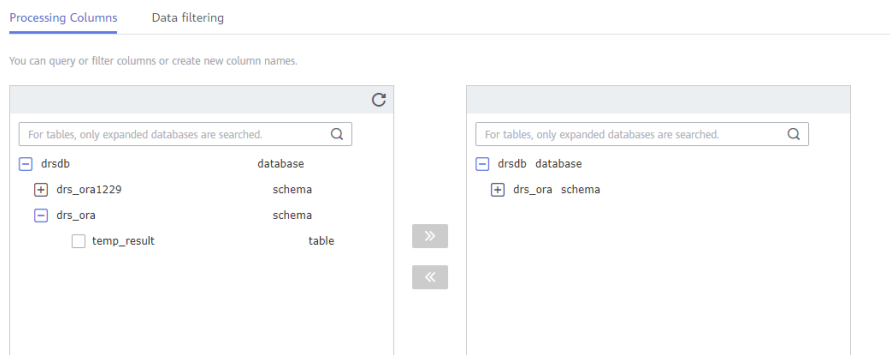
Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

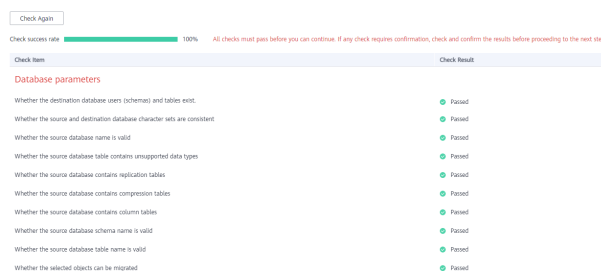
Figure 4-153 Processing data



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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-154 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-155 Task startup settings

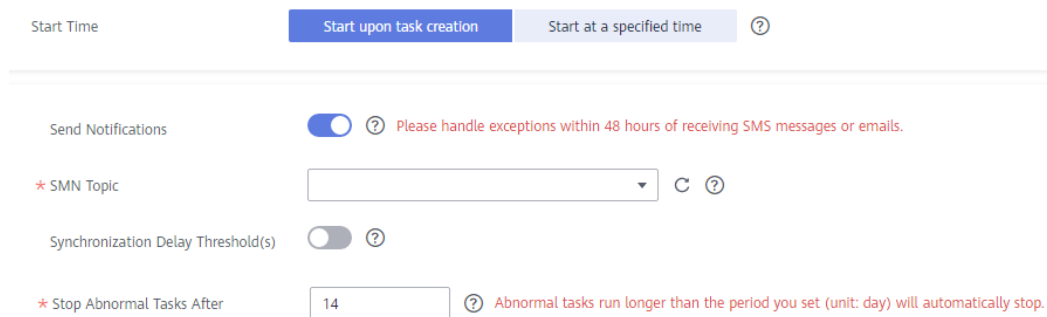



Table 4-176 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.18 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 4-177 Supported databases

Source DB	Destination DB
GaussDB primary/standby	Kafka 0.11 or later

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	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: INTEGER, TINYINT, SMALLINT, BIGINT, 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"> - 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - 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.

Table 4-179 Database user permission

Type	Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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.

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 [What Is the Impact of DRS on 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 4-180 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. ● Other notes: <ul style="list-style-type: none"> - The source must be a GaussDB primary/standby instance on Huawei Cloud. - 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.

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. • 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 <i>GaussDB Developer Guide</i>.
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> – When an 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 real-time incremental synchronization task, you may need to manually delete residual replication slots from the source database. For details, see.

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 4-156 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-6131 [Help icon]

Description: [Text area] 0/256 [Help icon]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Table 4-181 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-157 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 | DDM | GaussDB(for MySQL) Distributed Edition | GaussDB Distributed | **GaussDB Primary/Standby** | DDS | PostgreSQL | GaussDB(for MySQL)

* Destination DB Engine: Oracle | GaussDB(DWS) | GaussDB Distributed | GaussDB Primary/Standby | **Kafka**

* Network Type: Public network ⓘ
 DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.

* Instance Type: **Single** | Primary/Standby

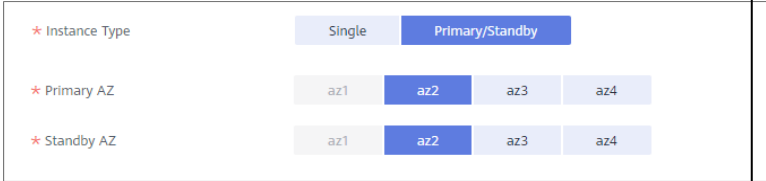
* Source DB Instance: No DB instance available ⓘ View DB Instance View Unselectable DB Instance

Synchronization Instance Subnet: Select the 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.

Table 4-182 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	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby. Note that if Primary/Standby is selected, data in tables without primary keys may be duplicate or inconsistent when a primary/standby switchover is performed.</p> <ul style="list-style-type: none">- Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization.- Single: The single-node deployment architecture is used, which is cost-effective.
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 Subnet 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 ensure that there are available IP addresses. 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. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p>

Parameter	Description
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-158 AZ</p> 

- DRS instance specifications

Figure 4-159 Specifications



Table 4-183 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-160 Enterprise projects and tags

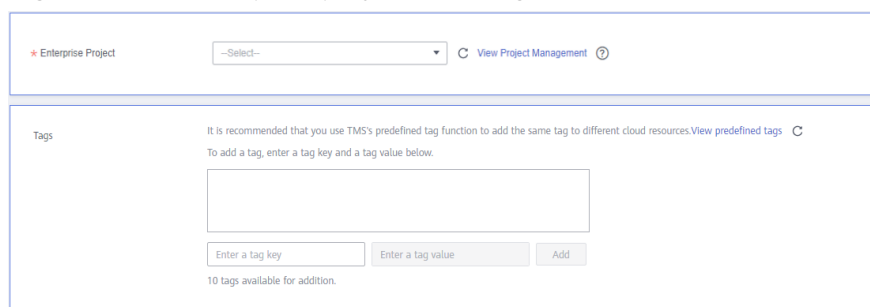


Table 4-184 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-161 Source database information

Source Database

DB Instance Name ads-ces-source

Database Username

Database Password

✔ Test successful

Table 4-185 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 4-162 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Connection Method

✔ Test successful

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.
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 4-163 Synchronization mode

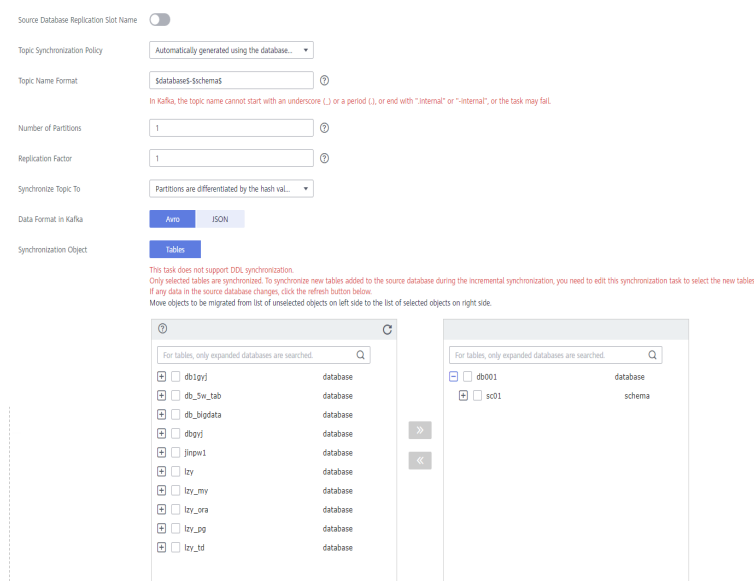



Table 4-187 Synchronization Object

Parameter	Description
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 .

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, 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-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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-164 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-165 Task startup settings

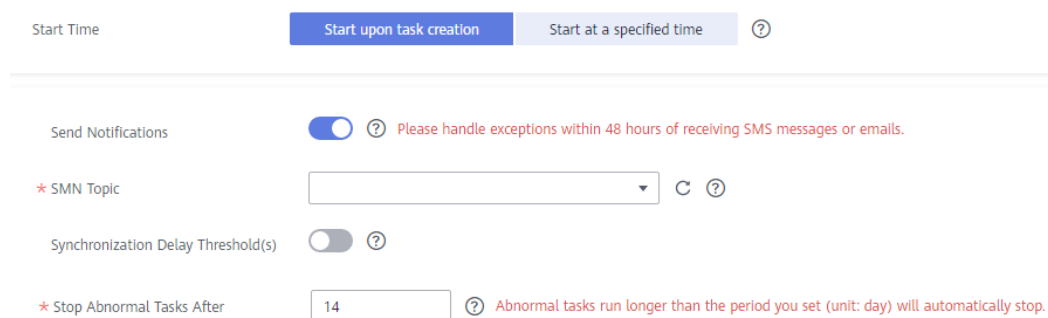



Table 4-188 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.19 From GaussDB Primary/Standby to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-189 Supported databases

Source DB	Destination DB
GaussDB primary/standby	GaussDB(DWS) cluster

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
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, 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 selected.- Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase.- 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.

Table 4-191 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	The user has the sysadmin role or the following permissions: CONNECT permission for DATABASE, USAGE permission for SCHEMA, and SELECT or UPDATE (required only for tables without primary keys) permission for tables.	The user has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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.
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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 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 /<.>\\'\ ?! - 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"> ● Other notes: <ul style="list-style-type: none"> - The source must be a GaussDB primary/standby instance on Huawei Cloud. - 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.
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.
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. ● 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 <i>GaussDB Developer Guide</i>.

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	<ul style="list-style-type: none">• Stop a task normally. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted.• Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby.

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 4-166 Synchronization task information




Table 4-192 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-167 Synchronization instance details

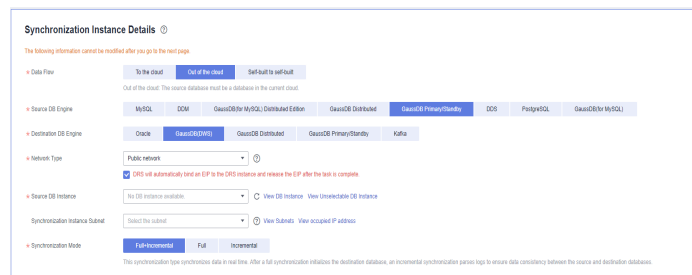


Table 4-193 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .

Parameter	Description
Destination DB Engine	Select GaussDB(DWS) .
Network Type	The public network is used as an example. Available options: Public network and VPN or Direct Connect
Source DB Instance	The GaussDB primary/standby instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full+Incremental , Full , and Incremental . Full+Incremental is used as an example. – 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.

- DRS instance specifications

Figure 4-168 Specifications



Table 4-194 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-169 Enterprise projects and tags




Table 4-195 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 4-170 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-196 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 4-171 Destination database information

Destination Database

IP Address or Domain Name 

Database Username

Database Password 

Ensure that the entered addresses belong to the same DB instance

Table 4-197 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 4-172 Synchronization mode

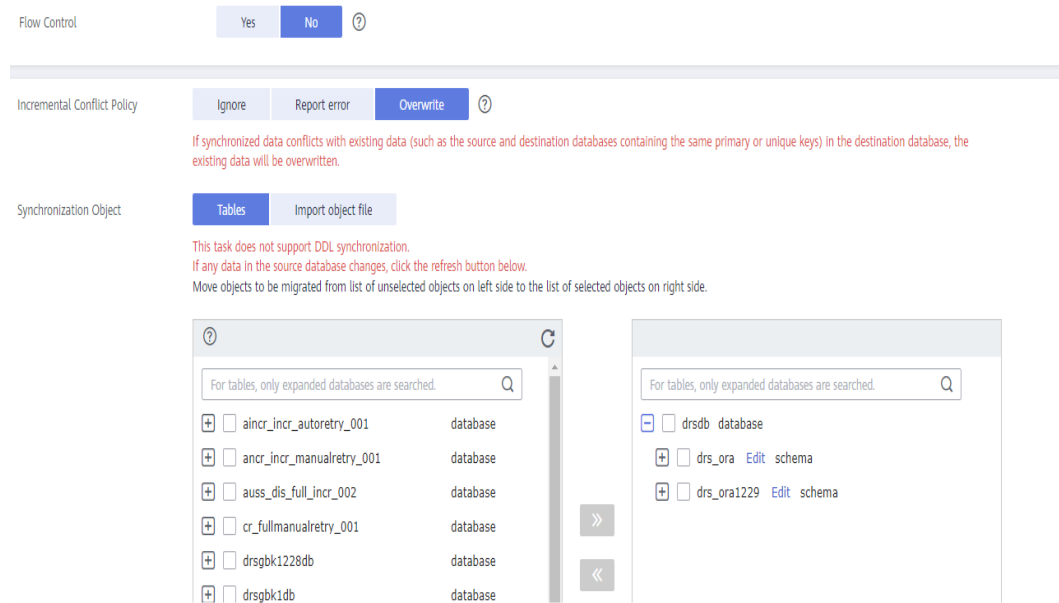
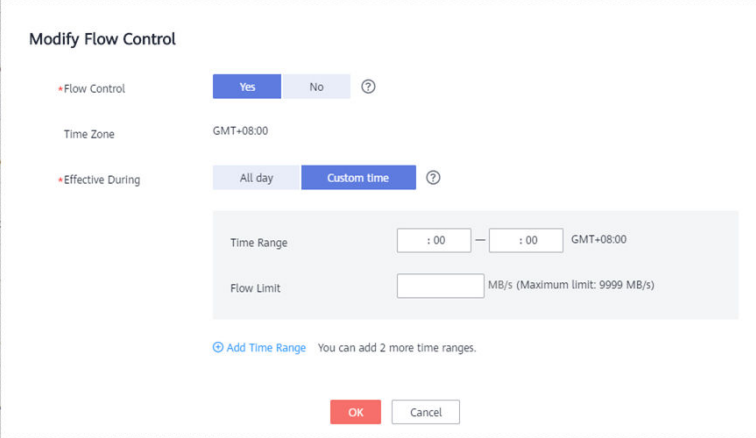



Table 4-198 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-173 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.


Figure 4-174 Parameter settings

Full Synchronization Settings

* Synchronization Object Type Table structure Data Constraint(excluding foreign keys)

* Concurrent Export Tasks   

* Concurrent Import Tasks   

* Rows per Shard 

Incremental Replay Settings




* Concurrent Replay Tasks   

Table 4-199 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none">• 0 All tables are not sharded, and each table is synchronized as a whole.• Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded.	520000

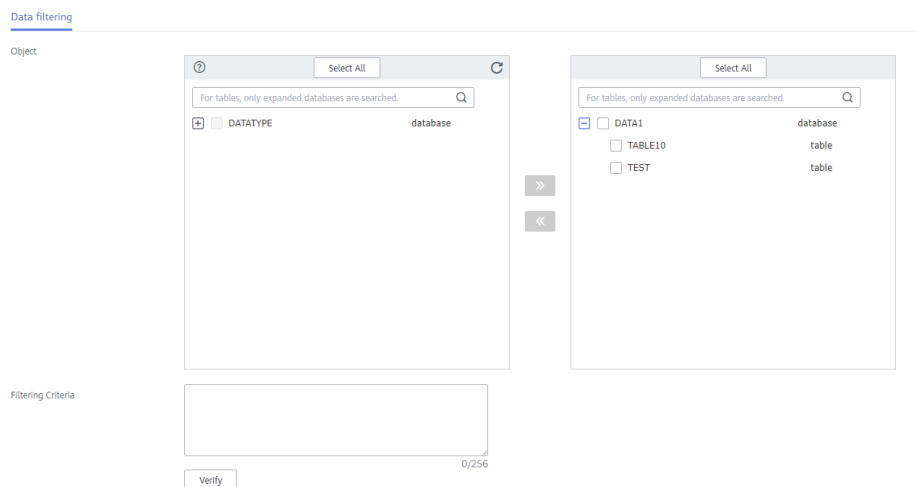
Table 4-200 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

Figure 4-175 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-176 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-177 Task startup settings

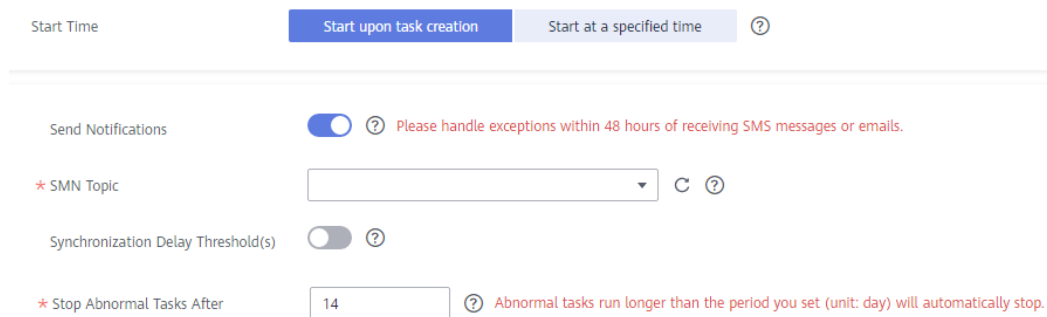



Table 4-201 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.20 From GaussDB Primary/Standby to GaussDB Distributed

Supported Source and Destination Databases

Table 4-202 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 4-203 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-203 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-204](#). 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.

Table 4-204 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. 	<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. 	<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.

Type	Full	Incremental	Full+Incremental
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, and INDEX permissions for tables. 	<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, and INDEX permissions for tables. 	<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, and INDEX permissions for 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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-205 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"> - The destination database must be a GaussDB distributed instance on the current cloud. - 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. - 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. - 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.
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.

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 <i>GaussDB Developer Guide</i>.
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.
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.

Type	Constraints
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby.

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 4-178 Synchronization task information

Table 4-206 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-179 Synchronization instance details

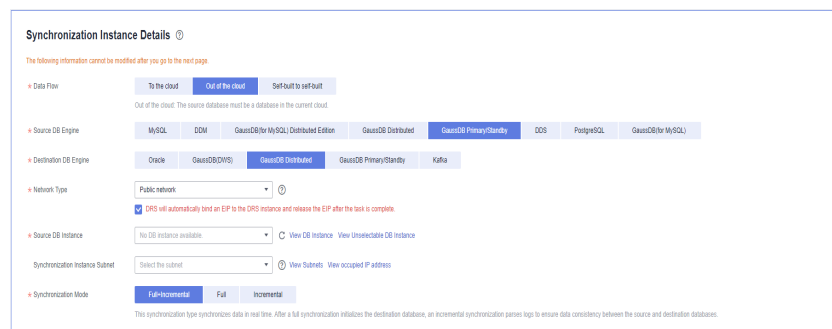


Table 4-207 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB Primary/Standby .
Destination DB Engine	Select GaussDB Distributed .
Network Type	The public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 4-180 Enterprise projects and tags

Table 4-208 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-181 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-209 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 4-182 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 4-210 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 4-183 Synchronization mode

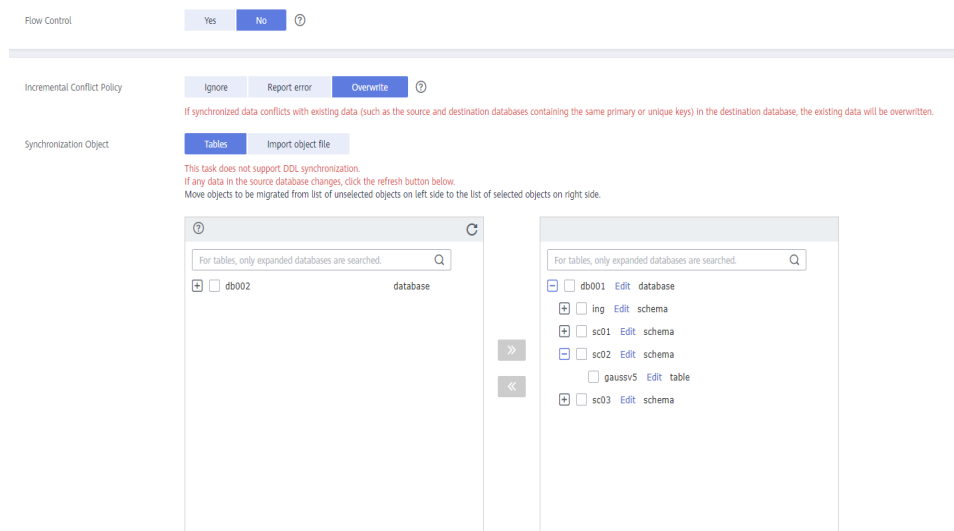
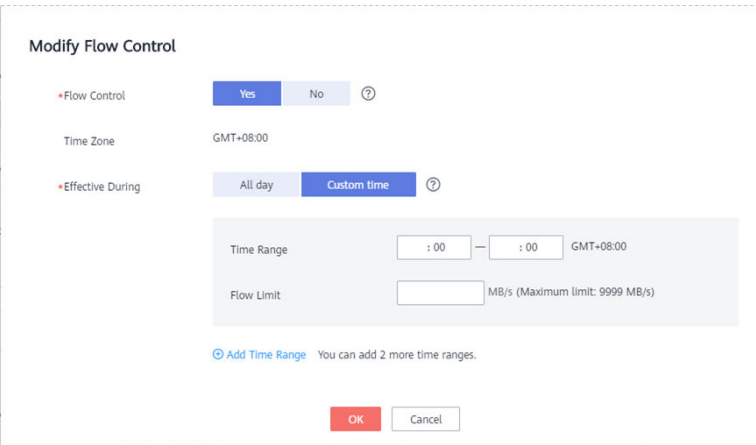



Table 4-211 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-184 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>Ignore and overwrite: Synchronization stability is prioritized, so tasks will not be interrupted as data conflicts occur.</p> <p>Report error: Data quality is prioritized. Any data conflicts are not allowed, so once a conflict occurs, the synchronization task fails and an error is reported. You need to manually find the cause of the fault. If the task is in the failed state for a long time, the storage space may be used up and the task cannot be restored.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in **Step 2** and click **Next**.

Figure 4-185 Parameter settings

Full Synchronization Settings

* Synchronization Object Type Table structure Data Constraint

* Concurrent Export Tasks 8

* Concurrent Import Tasks 8

* Rows per Shard

Incremental Replay Settings

* Concurrent Replay Tasks 64

Table 4-212 Full synchronization settings

Parameter	Description	Default Value
Synchronization Mode	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> ● 0 All tables are not sharded, and each table is synchronized as a whole. ● Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

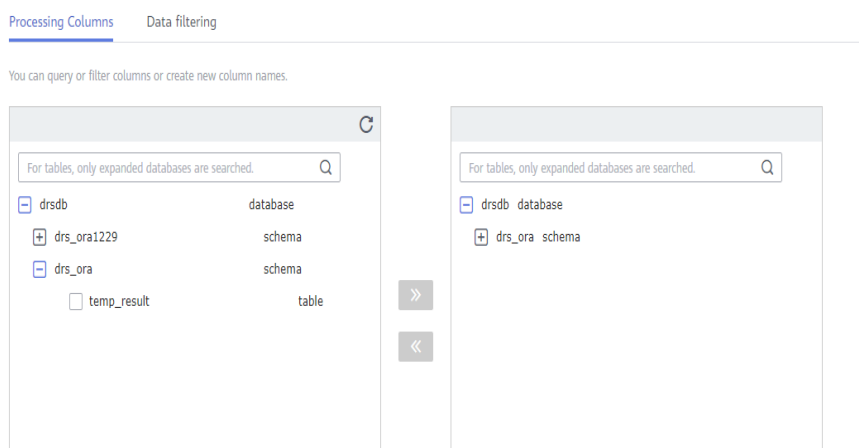
Table 4-213 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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**.

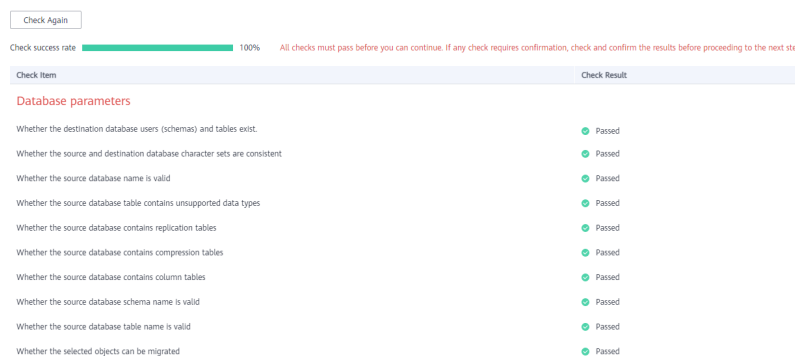
Figure 4-186 Processing data



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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-187 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-188 Task startup settings

The screenshot shows the 'Task startup settings' configuration page. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below this, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says 'Please handle exceptions within 48 hours of receiving SMS messages or emails.'
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field contains the value '14'. A red note says 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 4-214 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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
Synchronization Delay Threshold	<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 Notification.
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> <p>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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.21 From GaussDB Primary/Standby to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-215 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 4-216 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-216 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 4-217](#). 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.

Table 4-217 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. 	<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. 	<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.

Type	Full	Incremental	Full+Incremental
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, and INDEX permissions for tables. 	<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, and INDEX permissions for tables. 	<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, and INDEX permissions for 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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 4-218 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"> - The destination database must be a GaussDB primary/standby instance on the current cloud. - 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

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: ● SSL is not supported. ● 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. ● 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.
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.

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 <i>GaussDB Developer Guide</i>.
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.
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.

Type	Constraints
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally. After an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task. To forcibly stop an incremental or a full+incremental synchronization task, delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby.

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 4-189 Synchronization task information

Table 4-219 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-190 Synchronization instance details

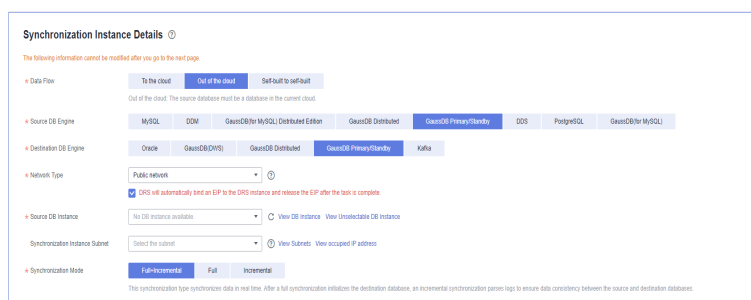


Table 4-220 Synchronization instance settings

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	The public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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.

- Enterprise Projects and Tags

Figure 4-191 Enterprise projects and tags

Table 4-221 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-192 Source database information

Source Database

DB Instance Name ads-ces-source

Database Username

Database Password

✔ Test successful

Table 4-222 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 4-193 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 4-223 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 4-194 Synchronization mode

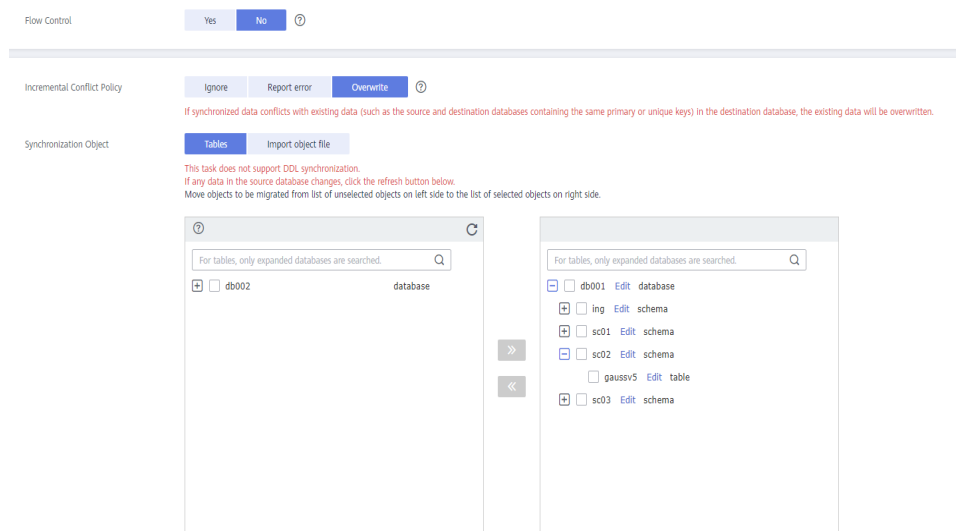
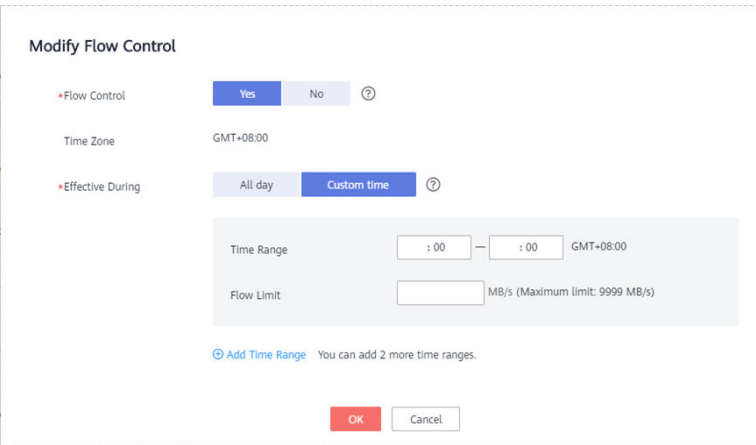



Table 4-224 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-195 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>Ignore and overwrite: Synchronization stability is prioritized, so tasks will not be interrupted as data conflicts occur.</p> <p>Report error: Data quality is prioritized. Any data conflicts are not allowed, so once a conflict occurs, the synchronization task fails and an error is reported. You need to manually find the cause of the fault. If the task is in the failed state for a long time, the storage space may be used up and the task cannot be restored.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in **Step 2** and click **Next**.

Figure 4-196 Parameter settings

Full Synchronization Settings

* Synchronization Object Type Table structure Data Constraint

* Concurrent Export Tasks 8

* Concurrent Import Tasks 8

* Rows per Shard

Incremental Replay Settings

* Concurrent Replay Tasks 64

Table 4-225 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Select whether to synchronize indexes based on the service requirements. Table structure and data are mandatory.	All options
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8

Parameter	Description	Default Value
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> • 0 All tables are not sharded, and each table is synchronized as a whole. • Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

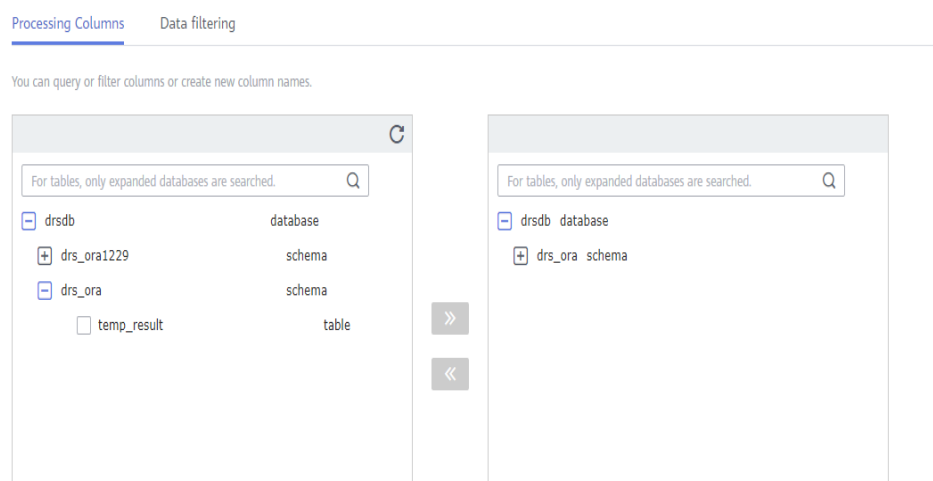
Table 4-226 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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 4-197 Processing data



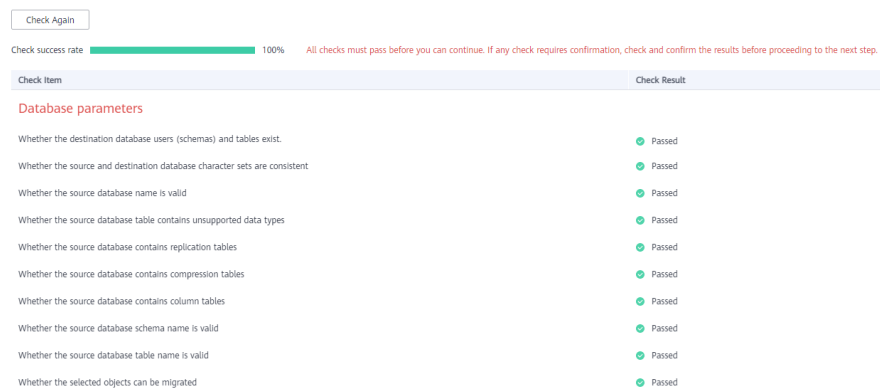
Step 7 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-198 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 8** On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-199 Task startup settings

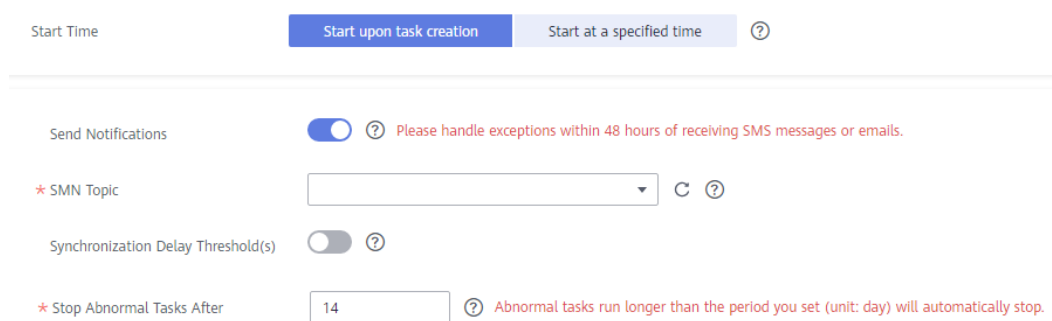



Table 4-227 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.22 From GaussDB(for MySQL) to MySQL

Supported Source and Destination Databases

Table 4-228 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">GaussDB(for MySQL)	<ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECSMySQL databases on other cloudsRDS for MySQL

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 4-229 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, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects. ● 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. ● Tables with storage engine different to MyISAM and InnoDB 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.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the 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. ● The source and destination databases cannot contain tables that have the same names but do not have primary keys. ● 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. ● Incremental synchronization supports some DDL operations. ● During table-level synchronization, the renamed table may conflict with other tables in the destination database. As a result, the task fails. Currently, renaming table is not supported. ● 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.

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 4-200 Synchronization task information

The screenshot shows a form with the following fields and options:

- Region:** A dropdown menu with a search icon.
- Project:** A dropdown menu.
- Task Name:** A text input field containing "DRS-B131".
- Description:** A text area with a character count of "0/256".

At the top, there is a warning message: "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."

Table 4-230 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-201 Synchronization instance details

The screenshot shows the 'Synchronization Instance Details' configuration page with the following settings:

- Data Flow:** Radio buttons for 'To the cloud', 'Out of the cloud', and 'Self-built to self-built'. 'Out of the cloud' is selected.
- Source DB Engine:** Tabs for MySQL, DDM, GaussDB(for MySQL) Distributed Edition, GaussDB Distributed, GaussDB Primary/Standby, DDS, PostgreSQL, and GaussDB(for MySQL). 'GaussDB(for MySQL)' is selected.
- Destination DB Engine:** Tabs for MySQL, Oracle, GaussDB(DWS), CSSES, and Kafka. 'MySQL' is selected.
- Network Type:** A dropdown menu set to 'Public network'.
- DRS Settings:** A checkbox 'DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.' is checked.
- Source DB Instance:** A dropdown menu with 'View DB Instance' and 'View Unselectable DB Instance' links.
- Synchronization Instance Subnet:** A dropdown menu with 'View Subnets' and 'View occupied IP address' links.
- Synchronization Mode:** Radio buttons for 'Full-consistent' and 'Incremental'. 'Full-consistent' is selected.

A note at the bottom states: "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 4-231 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select MySQL .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-202 Specifications



Table 4-232 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-203 Enterprise projects and tags




Table 4-233 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 4-204 Source database information

Source Database

DB Instance Name ads-000-source

Database Username root

Database Password

Test Connection
✔ Test successful

Table 4-234 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 4-205 Destination database information

Destination Database

IP Address or Domain Name [blurred]

Port [blurred]

Database Username [blurred]

Database Password

SSL Connection

Test Connection
✔ Test successful

Table 4-235 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 4-206 Synchronization mode

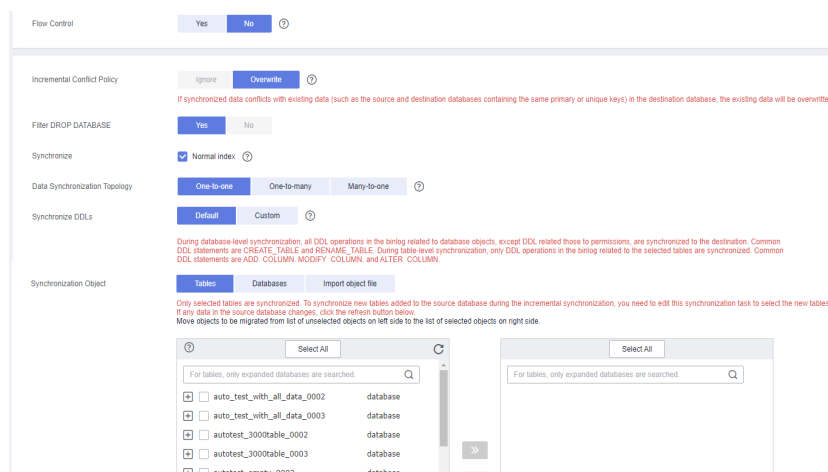



Table 4-236 Synchronization mode and object

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.● 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.
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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</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. <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 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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](#).

Figure 4-207 Processing data



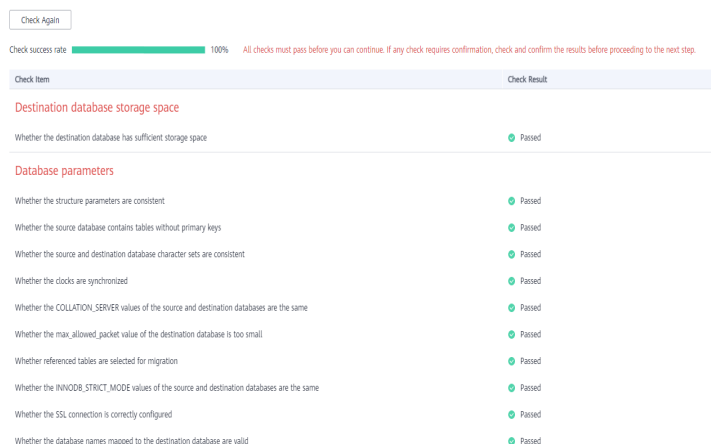
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-208 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-209 Task startup settings





Table 4-237 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.23 From GaussDB(for MySQL) to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-238 Supported databases

Source DB	Destination DB
GaussDB(for MySQL)	GaussDB(DWS) cluster

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 4-239 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, CREATE, and REFERENCES.
Synchronization object	<ul style="list-style-type: none"> • Only tables, indexes, and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraint, and unique index, 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. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the 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. ● Indexes can be synchronized, regardless of index types. After the synchronization, the index type becomes the default type. ● 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. ● 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. ● 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: <ul style="list-style-type: none"> – CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, ALTER_COLUMN, DROP_CONSTRAINT, ADD_CONSTRAINT, CREATE_INDEX, DROP_INDEX and RENAME_INDEX are supported. – DROP_DATABASE, DROP_TABLE, TRUNCATE_TABLE, CREATE_VIEW and DROP_VIEW are not supported. – Tables can be created. For example: <pre style="margin: 0;">create table `ddl_test` (id int, c1 varchar(25), primary key(id)); create table `ddl_test_gho` like `ddl_test`;</pre> – Columns in a table can be added and modified, but cannot be deleted. For example:

Type	Restrictions
	<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> <ul style="list-style-type: none"> - 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. If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. For example: <pre>alter table `ddl_test` modify column `c1` varchar(50);</pre> - During table-level synchronization, the renamed table may conflict with other tables in the destination database. As a result, the task fails. Currently, renaming table is not supported. - 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. • 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. • 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 4-210 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-8131 [Help icon]

Description: [Text area] [Help icon]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Table 4-240 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-211 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 for MySQL Distributed Edition | GaussDB Distributed | GaussDB Primary/Standby | DDS | PostgreSQL | **GaussDB for MySQL**
- Destination DB Engine:** MySQL | Oracle | **GaussDB(DWS)** | CSSES | Kafka
- Network Type:** Public network
- Source DB Instance:** [Dropdown menu] | View DB Instance | View Unselectable DB Instance
- Synchronization Instance Subnet:** [Dropdown menu] | View Subnets | View occupied IP address
- Synchronization Mode:** **Full-Incremental**

Out of the cloud. The source database must be a database in the current cloud.

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.

DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.

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 4-241 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .

Parameter	Description
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select GaussDB(DWS) .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source DB Instance	An available GaussDB(for MySQL) instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	<p>– 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> <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 instance specifications

Figure 4-212 Specifications



Table 4-242 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-213 Enterprise projects and tags

Table 4-243 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 4-214 Source database information

Source Database

DB Instance Name rds-ow-source

Database Username

Database Password

✔ Test successful

Table 4-244 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 4-215 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	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 select **Data Synchronization Topology** and synchronize incremental DDL operations. To use these functions, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Figure 4-216 Synchronization mode

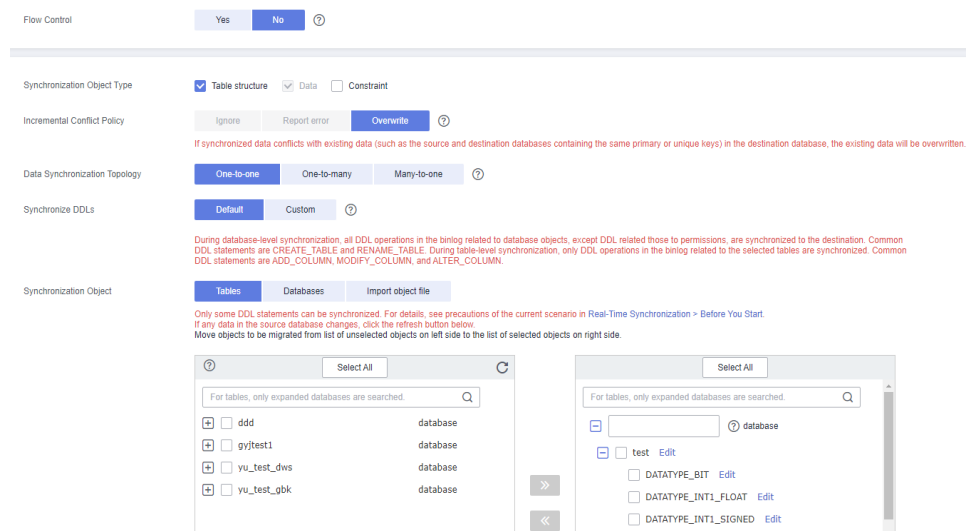
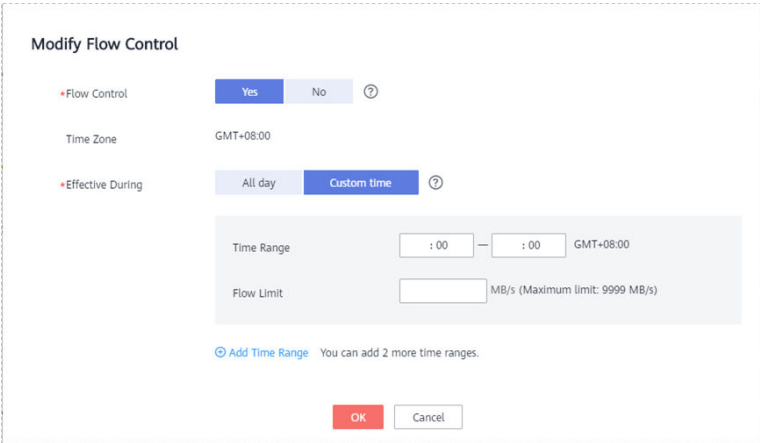



Table 4-246 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-217 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>Available options: Table structure, Data, and Constraint</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>
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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</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. <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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 4-218 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
test_0726.common_1	test_0726.common_1	--	--	--	Add
test_0726.common	test_0726.common	--	--	--	Add

Batch Apply 

 **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.

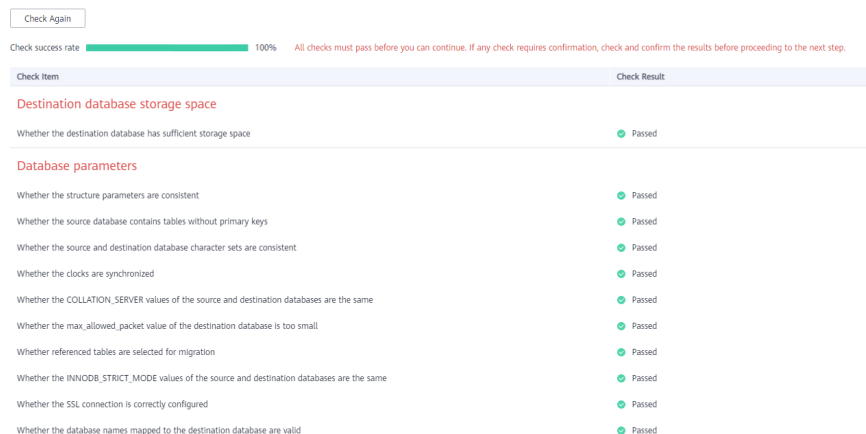
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-219 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-247 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.24 From GaussDB(for MySQL) to Kafka

Supported Source and Destination Databases

Table 4-248 Supported databases

Source DB	Destination DB
GaussDB(for MySQL)	Kafka

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](#).

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 [What Is the Impact of DRS on 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 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.

Table 4-249 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, REPLICATION CLIENT, and RELOAD.
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 can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.
Source database	<ul style="list-style-type: none"> • The source database must be the primary node of the 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.
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.) • 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 4-221 Synchronization task information




Table 4-250 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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

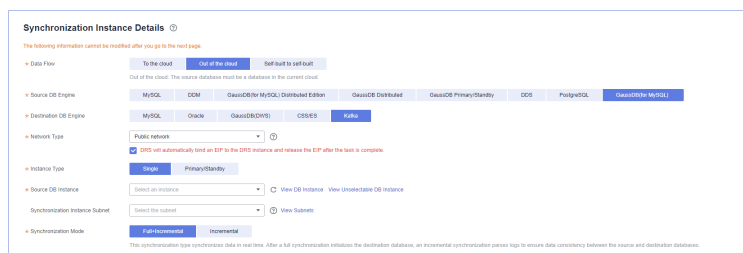
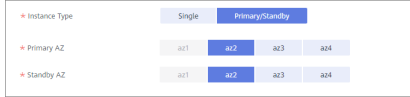


Table 4-251 Synchronization instance settings

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

Parameter	Description
Destination DB Engine	Select Kafka .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none"> - Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. - Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</p>
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 Subnet 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 ensure that there are available IP addresses. 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.

Parameter	Description
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-223 AZ</p> 

- DRS instance specifications

Figure 4-224 Specifications



Table 4-252 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-225 Enterprise projects and tags

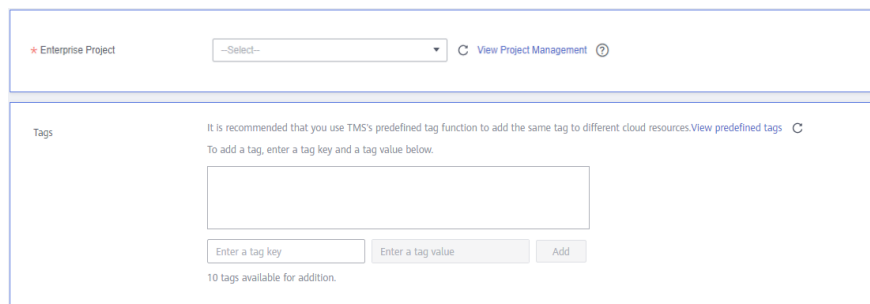


Table 4-253 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-226 Source database information

Source Database

DB Instance Name rds-omw-source

Database Username

Database Password

 Test successful

Table 4-254 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.


Parameter	Description
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 4-227 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.


Security Protocol 

Table 4-255 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 4-228 Synchronization mode

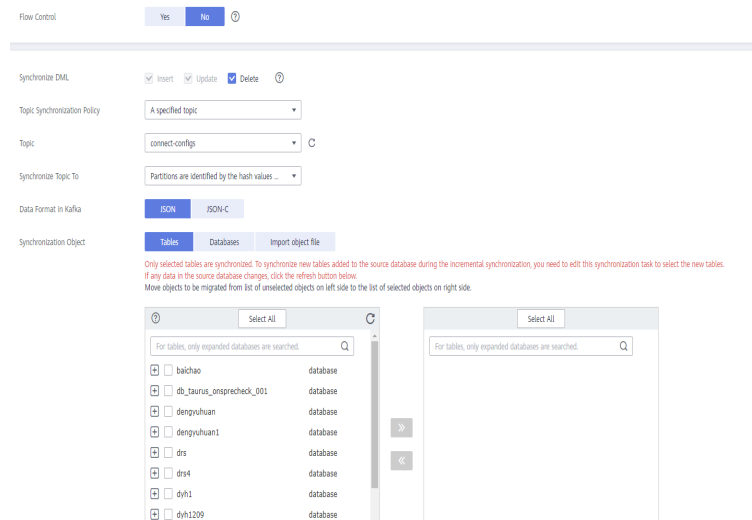
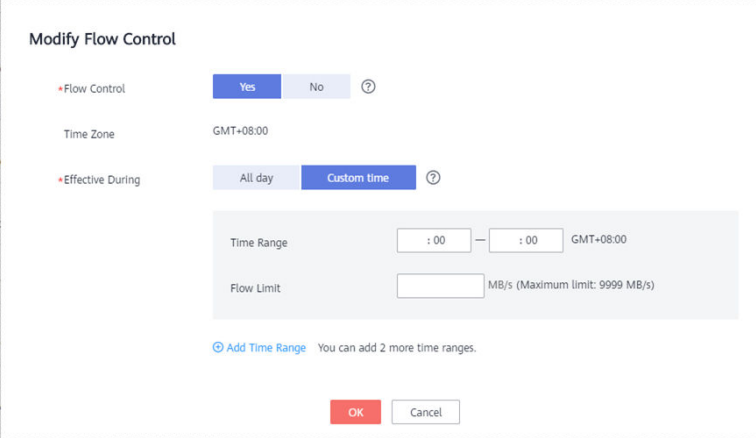



Table 4-256 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-229 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.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</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 start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</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	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 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.

Parameter	Description
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 when Synchronization Mode is set to Incremental in Step 2.● 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>
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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. This function is available when **Synchronization Mode** is set to **Incremental** in [Step 2](#).

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 4-230 Processing data



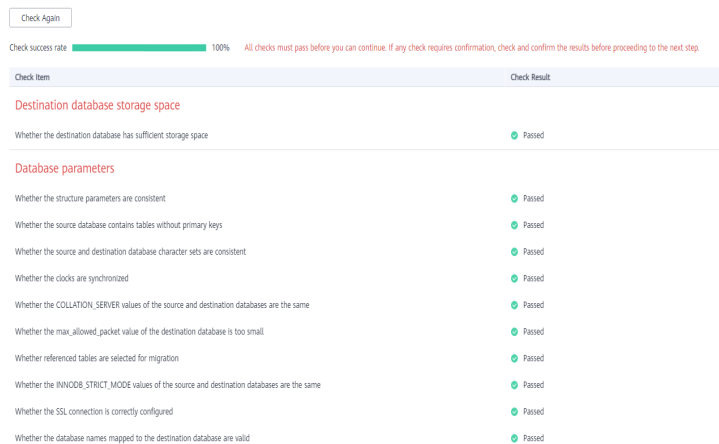
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-231 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-232 Task startup settings

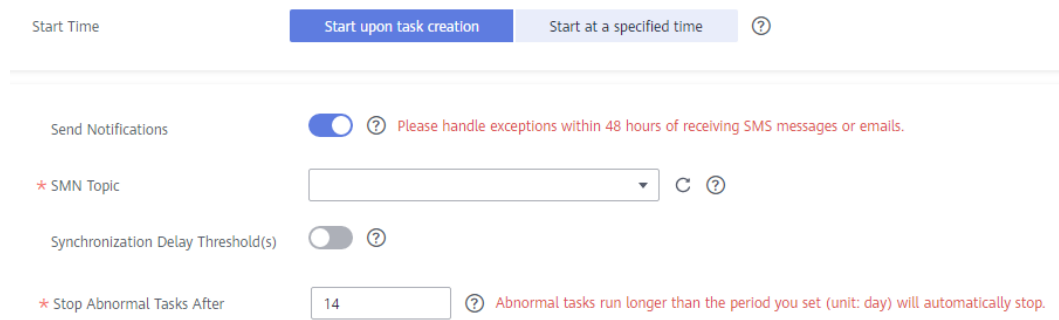



Table 4-257 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.25 From GaussDB(for MySQL) to CSS/ES

Supported Source and Destination Databases

Table 4-258 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • GaussDB(for MySQL) 	<ul style="list-style-type: none"> • ElasticSearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9 and 7.10

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 4-259 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
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.

Type	Restrictions
Source database	<ul style="list-style-type: none">• The source database must be the primary node of the 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 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.● 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.● During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail.● Incremental synchronization filters out all DDL operations.

Type	Restrictions
	<ul style="list-style-type: none"> • 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 4-233 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: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: [Text input: DRS-6131]

Description: [Text area]

Table 4-260 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-234 Synchronization instance details

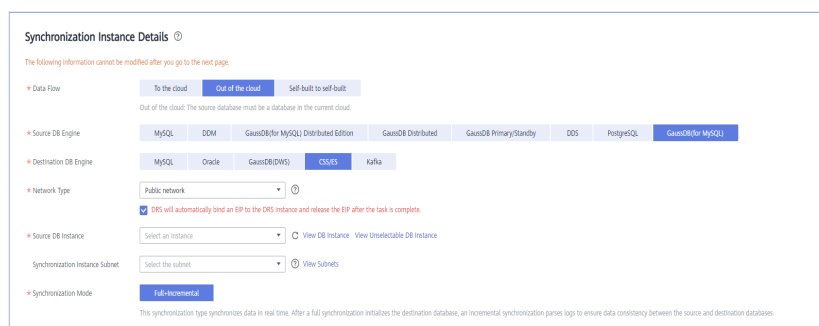


Table 4-261 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select GaussDB(for MySQL) .
Destination DB Engine	Select CSS/ES .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
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.

- DRS instance specifications

Figure 4-235 Specifications



Table 4-262 Specifications

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications .

- Enterprise Projects and Tags

Figure 4-236 Enterprise projects and tags

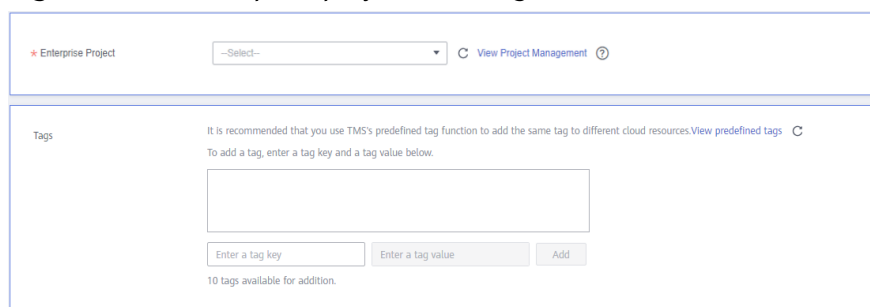


Table 4-263 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 4-237 Source database information

Source Database

DB Instance Name ods-oms-source

Database Username

Database Password

✔ Test successful

Table 4-264 Source database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance you selected when creating the task. This parameter cannot be changed.

Parameter	Description
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 4-238 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

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

Encryption Certificate

Table 4-265 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.

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 not enabled, 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 4-239 Synchronization mode

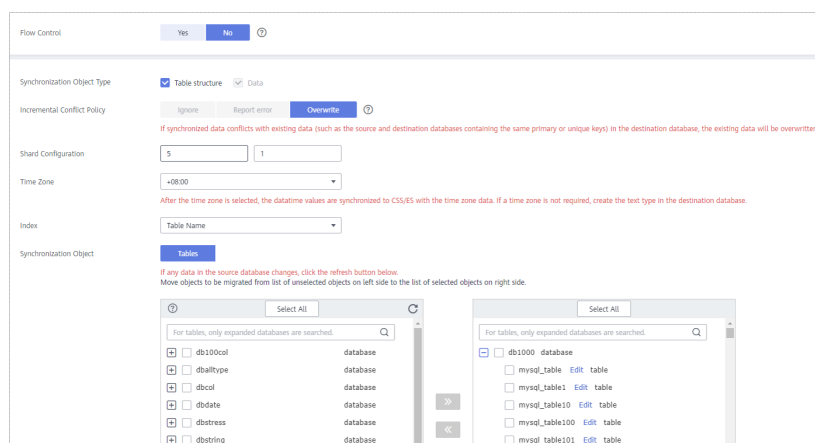
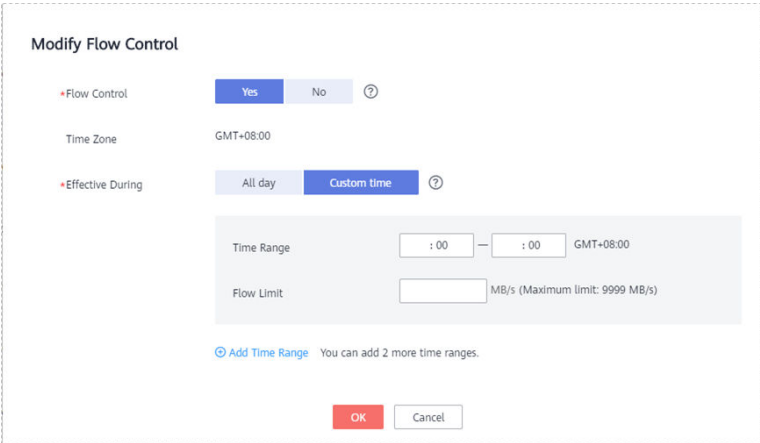



Table 4-266 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-240 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>Available options: Table structure and Data</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>
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 the time zone is selected, the data of the datetime type will be synchronized to CSS/ES with the selected time zone. If the time zone is not required, you need to manually create the document type on the target end.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces.

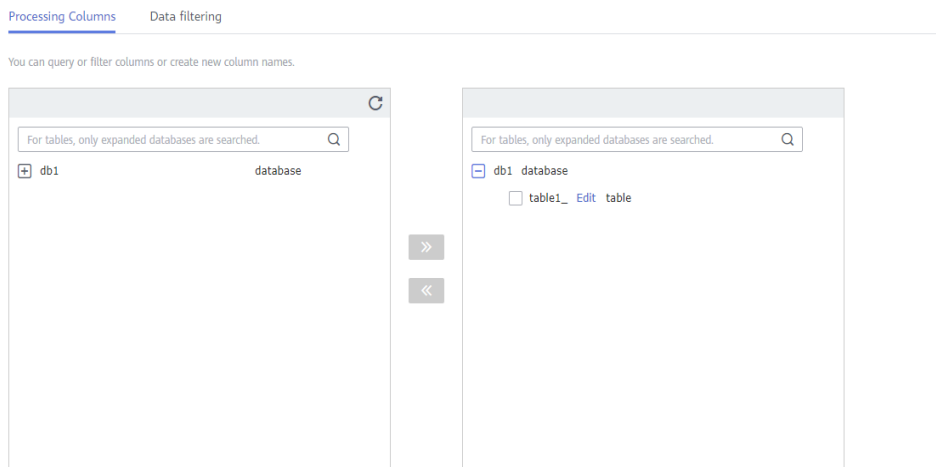
Step 5 On the **Process Data** page, set the filtering rules for data processing.

- 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.

Figure 4-241 Processing data



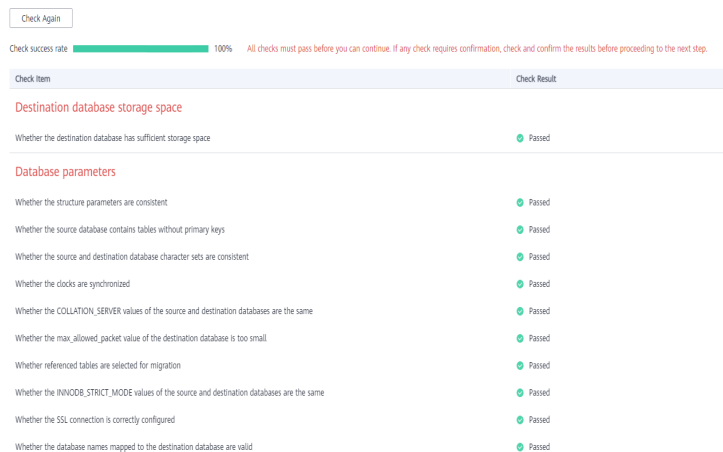
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-242 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-243 Task startup settings





Table 4-267 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.26 From GaussDB(for MySQL) to Oracle

Supported Source and Destination Databases

Table 4-268 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • GaussDB(for MySQL) instances 	<ul style="list-style-type: none"> • On-premises Oracle databases • Oracle databases on an ECS

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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

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 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.

Table 4-269 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.
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.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● The source database must be the primary node of the 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. ● 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. ● 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. ● 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.

Type	Restrictions
	<ul style="list-style-type: none"> • 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 4-244 Synchronization task information

Table 4-270 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 4-245 Synchronization instance details

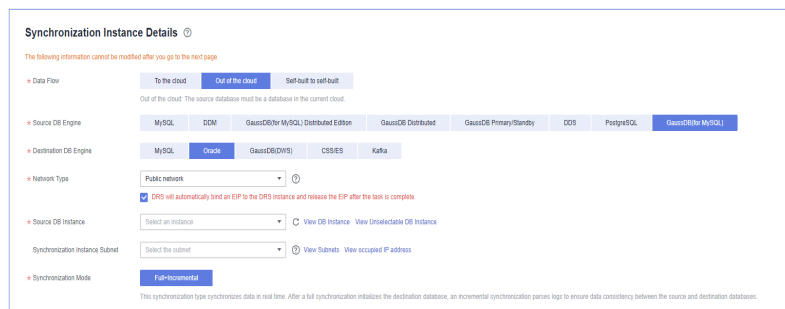


Table 4-271 Synchronization instance settings

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: Public network, VPC, VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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.

- DRS instance specifications

Figure 4-246 Specifications



Table 4-272 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 4-247 Enterprise projects and tags

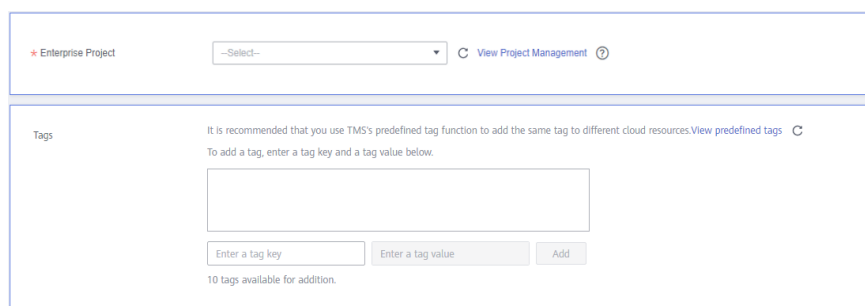


Table 4-273 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> – This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. – 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 ends.

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 4-248 Source database information

Source Database

DB Instance Name rds-cn-sou

Database Username

Database Password

✔ Test successful

Table 4-274 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 4-249 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 4-275 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 not enabled, 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 4-250 Synchronization mode

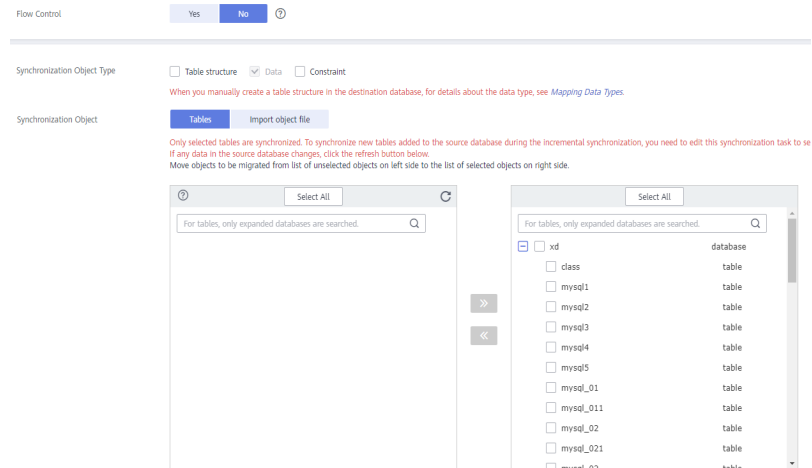
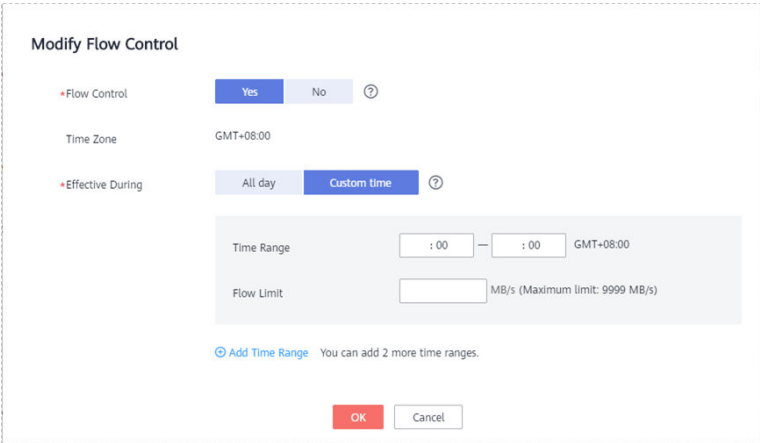



Table 4-276 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 4-251 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>Available options: Table structure, Data, and Constraint</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.
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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

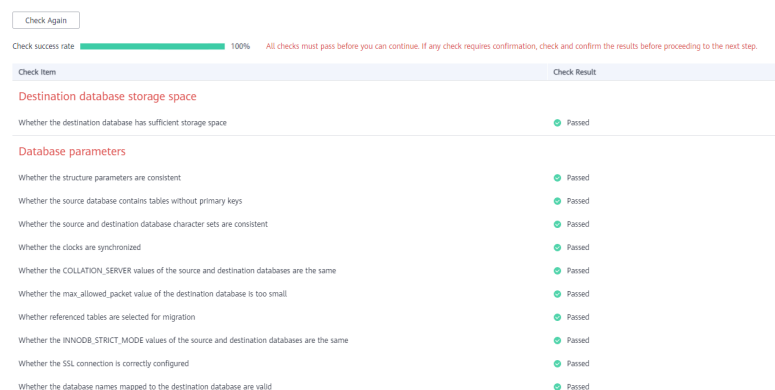
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 4-252 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 4-253 Task startup settings

The screenshot shows the 'Task startup settings' configuration page. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below this, there are four main settings:


- Send Notifications:** A toggle switch is turned on. A red note below it says: 'Please handle exceptions within 48 hours of receiving SMS messages or emails.'
- SMN Topic:** A dropdown menu is empty. To its right are refresh and help icons.
- Synchronization Delay Threshold(s):** A toggle switch is turned off. A help icon is to its right.
- 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 4-277 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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
Synchronization Delay Threshold	<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 Notification.
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> <p>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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----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

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](#).

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 [What Is the Impact of DRS on 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 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.

Table 5-2 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, REPLICATION CLIENT, and RELOAD.
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 can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none"> • 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. • The database and table names in the source database cannot contain non-ASCII characters, or special characters '<>\'\" • 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.

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.) ● 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. ● 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: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 [Help icon]

Description: [Text area] 0/256 [Help icon]

Table 5-3 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 [Help icon]

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 | PostgreSQL

* Destination DB Engine: MySQL | CSSES | **Kafka**

* Network Type: Public network [Help icon]

DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.

* Instance Type: **Single** | Primary-Standby

* VPC: default [View VPC]

Synchronization Instance Subnet: default_subnet [Help icon] [View Subnets] [View occupied IP address]

* Security Group: default [View]

* Synchronization Mode: **Full-Incremental** | Incremental

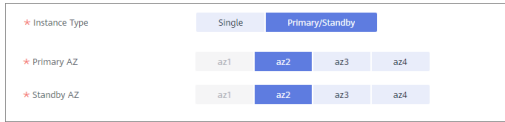
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases.

* AZ: **az1** | az2 | az3

AZ where the DRS instance is created. Selecting an AZ where the source or destination database is located provides better performance.

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
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none">- Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization.- Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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 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. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.
AZ	<p>Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-3 AZ</p> 

- DRS instance specifications

Figure 5-4 Specifications



Table 5-5 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 5-5 Enterprise projects and tags

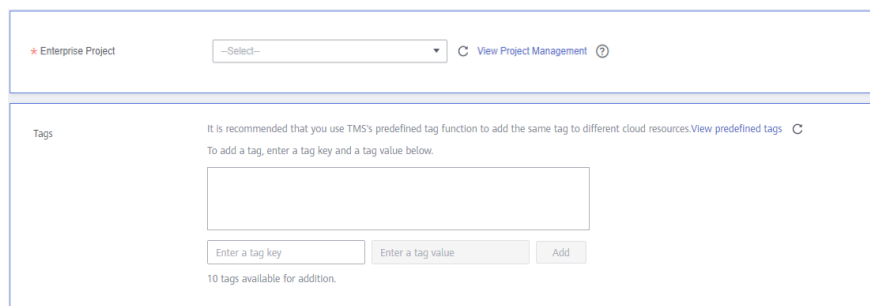


Table 5-6 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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, 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 ?
Ensure that the entered addresses belong to the same DB instance.

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 Select

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

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.
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 not enabled, 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 or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

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**.

Figure 5-8 Synchronization mode

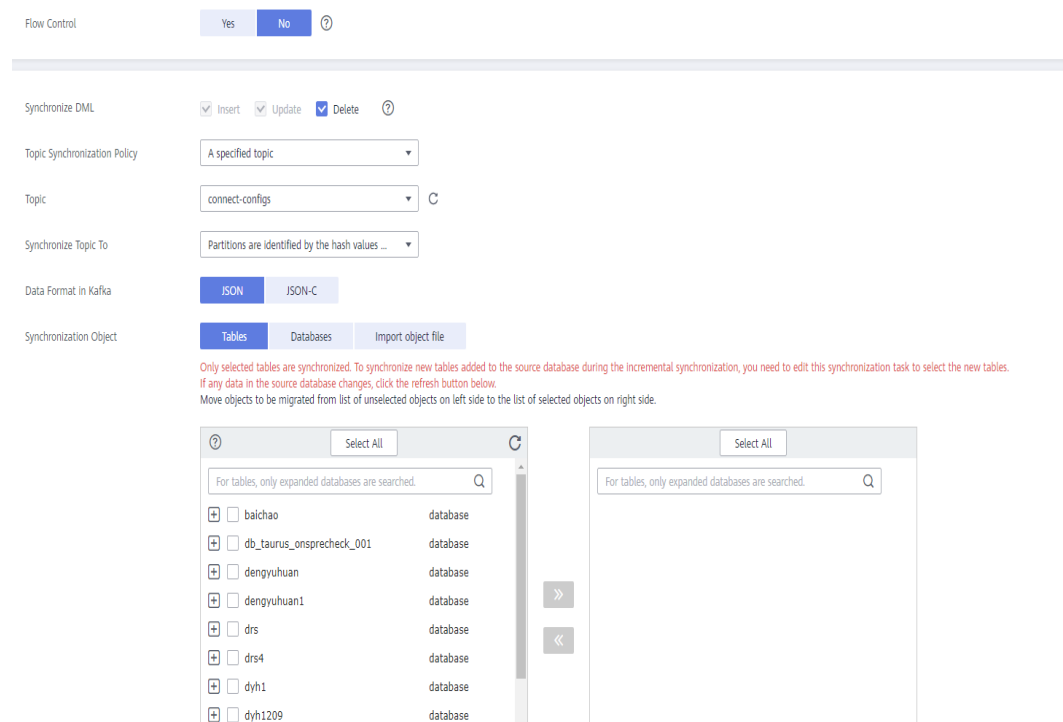
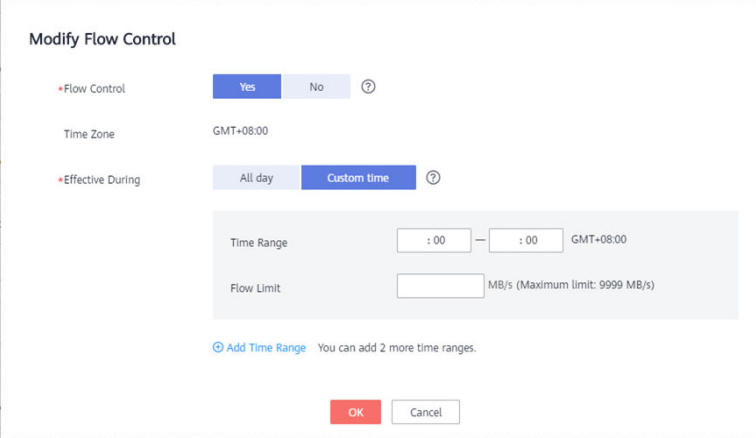



Table 5-9 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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.
Synchronize DML	<p>Select the DML operations to be synchronized. By default, all DML operations are selected.</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 start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</p>

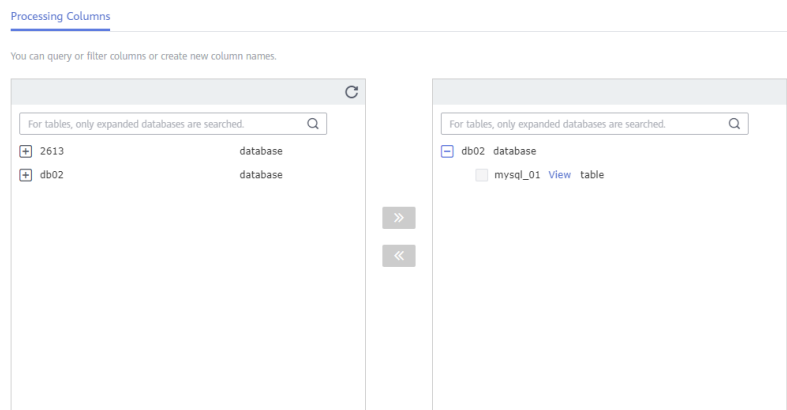
Parameter	Description
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	<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>
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.

Parameter	Description
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 when Synchronization Mode is set to Incremental in Step 2. ● 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>
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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. This function is available when **Synchronization Mode** is set to **Incremental** in [Step 2](#).

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 5-10 Processing data



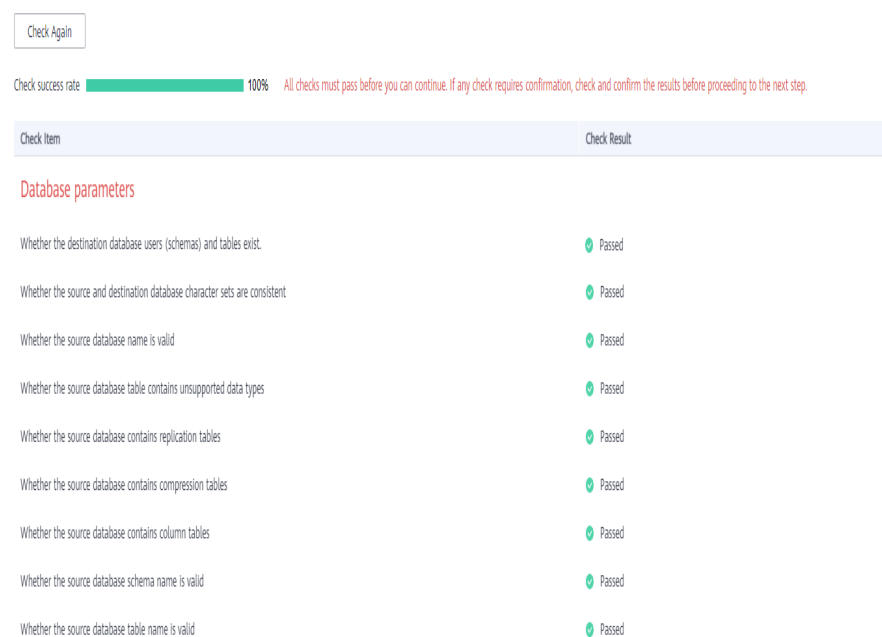
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-11 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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

The screenshot shows the 'Task startup settings' interface. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below the tabs, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says: "Please handle exceptions within 48 hours of receiving SMS messages or emails."
- * SMN Topic:** A dropdown menu is empty. There are refresh and help icons.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- * Stop Abnormal Tasks After:** A text input field contains the number '14'. A red note says: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Table 5-10 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----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

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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.

Table 5-12 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.

Type	Restrictions
Source database	<ul style="list-style-type: none"> ● 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 ibttmp1. ● 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. ● 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. ● 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.

Type	Restrictions
Destination database	<ul style="list-style-type: none"><li data-bbox="587 297 1241 331">• The destination DB instance is running properly.<li data-bbox="587 338 1358 409">• 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. ● 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 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

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message: "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." Below the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "* Task Name" (a text input field containing "DRS-6131"), and "Description" (a text area with a "0/256" character count). Each of the last three fields has a help icon (a circle with a question mark) to its right.

Table 5-13 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 Subnet 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 ensure that there are available IP addresses. 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	Available options: Full+Incremental and Incremental – 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.
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.

- DRS instance specifications

Figure 5-15 Specifications



Table 5-15 Specifications

Parameter	Description
Specifications	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications .

- Enterprise Projects and Tags

Figure 5-16 Enterprise projects and tags

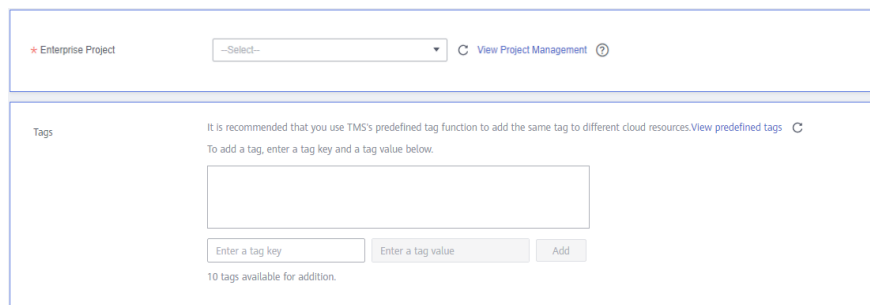


Table 5-16 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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, 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 ⓘ
Ensure that the entered addresses belong to the same DB instance.

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 5-17 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.
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 not enabled, 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

Database Username

Database Password

SSL Connection

Test successful

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 not enabled, 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

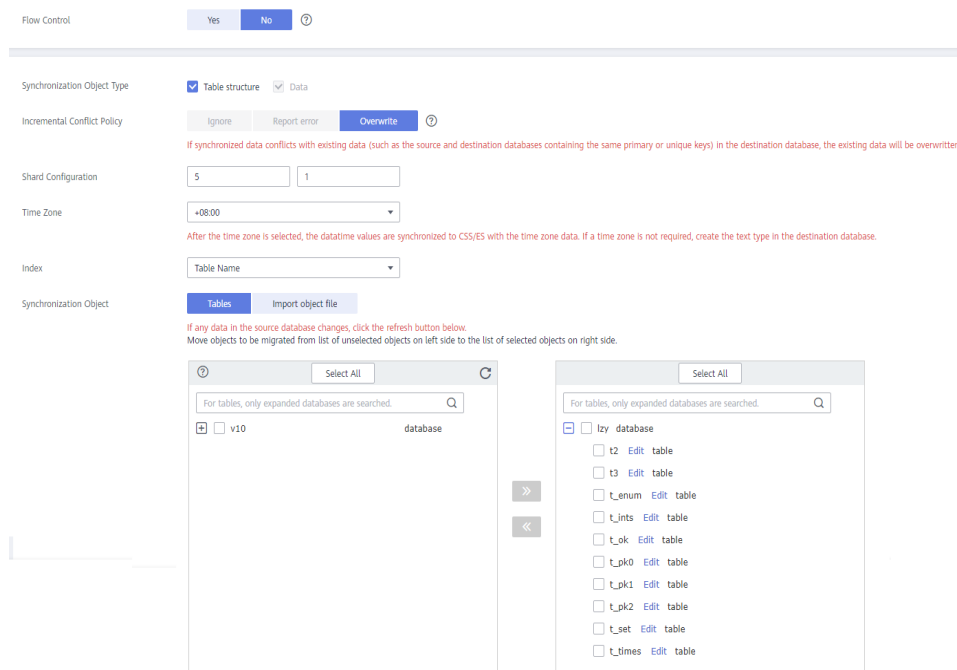
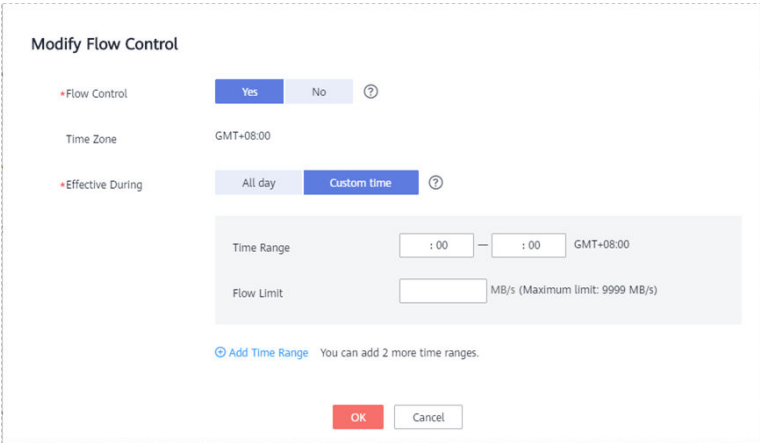



Table 5-19 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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.

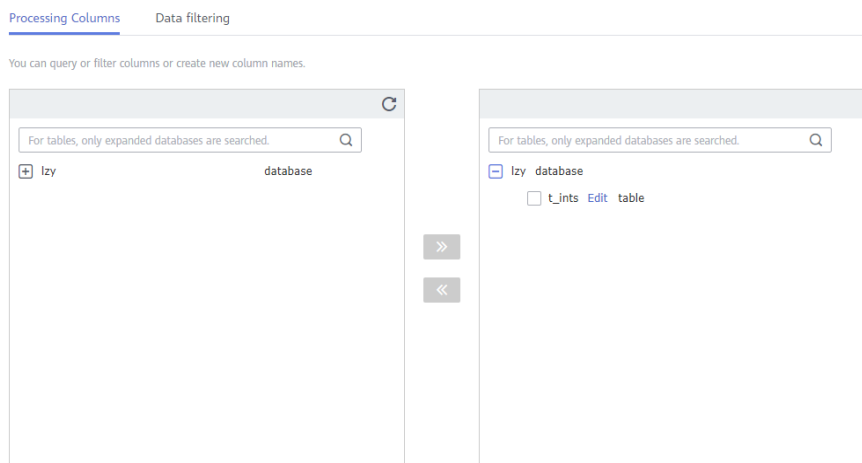
Parameter	Description
Synchronization Object Type	<p>Available options: Table structure and Data</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • 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>
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 the time zone is selected, the data of the datetime type will be synchronized to CSS/ES with the selected time zone. If the time zone is not required, you need to manually create the document type on the target end.</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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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



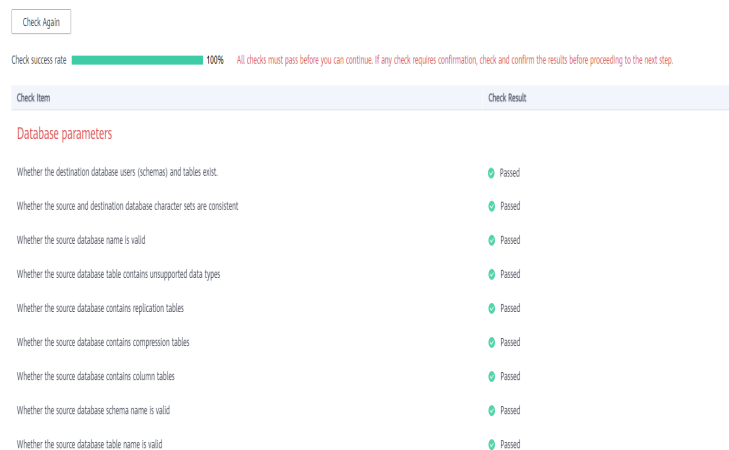
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-22 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-23 Task startup settings

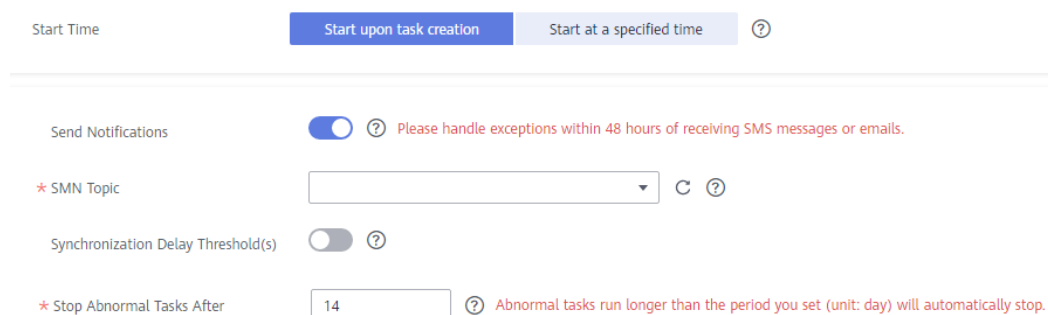



Table 5-20 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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>
Synchronization Delay Threshold	<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 Notification.
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 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, and the task

status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.3 From Oracle to Kafka

Supported Source and Destination Databases

Table 5-21 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises Oracle databases• Oracle databases on an ECS	<ul style="list-style-type: none">• Kafka

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](#).

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 [What Is the Impact of DRS on 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 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.

Table 5-22 Environment Constraints

Type	Constraint
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). - 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. - For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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: NTERVAL_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, 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 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. ● Tables whose default values contain expressions of functions cannot be synchronized. ● 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.
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.

Type	Constraint
Destination database	<ul style="list-style-type: none"><li data-bbox="587 297 1214 331">• The destination database is a Kafka database.

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. ● 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. ● 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 an incremental synchronization of table-level objects, renaming tables is not recommended.

Type	Constraint
	<ul style="list-style-type: none"> • 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.

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 5-24 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. X

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 ⓘ

Description: [Text area] ⓘ

0/256

Table 5-23 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-25 Synchronization instance details

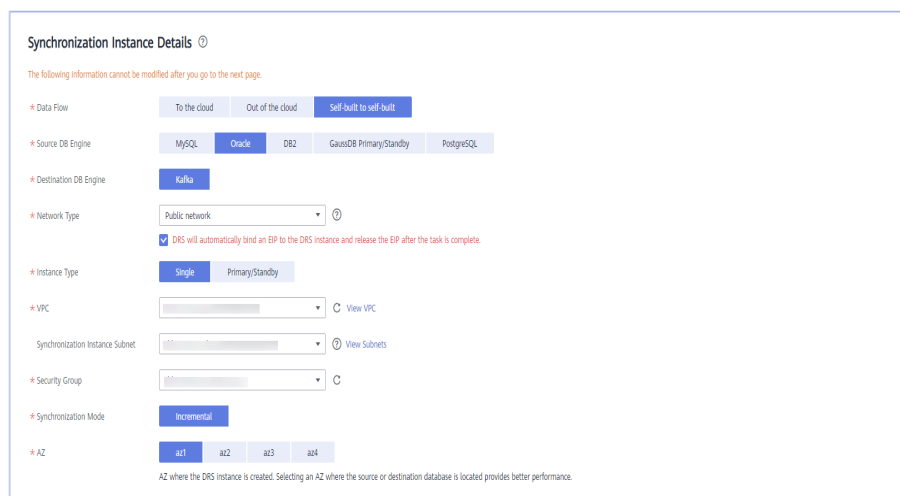



Table 5-24 Synchronization instance settings

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

Parameter	Description
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none"> – Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. – Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</p>
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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	<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>
AZ	<p>Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-26 AZ</p>  <p>The screenshot shows a configuration panel with three rows: <ul style="list-style-type: none"> Instance Type: Two radio buttons, 'Single' (unselected) and 'Primary/Standby' (selected). Primary AZ: Four radio buttons, 'az1' (unselected), 'az2' (selected), 'az3' (unselected), and 'az4' (unselected). Standby AZ: Four radio buttons, 'az1' (unselected), 'az2' (selected), 'az3' (unselected), and 'az4' (unselected). </p>

- DRS instance specifications

Figure 5-27 Specifications



Table 5-25 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 5-28 Enterprise projects and tags

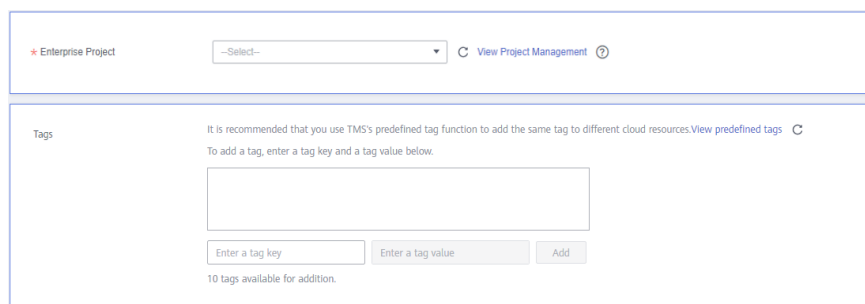


Table 5-26 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 5-29 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 5-27 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 not enabled, 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 5-30 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.


Security Protocol 

Table 5-28 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 a topic and objects to be synchronized, and then click **Next**.

Figure 5-31 Synchronization mode

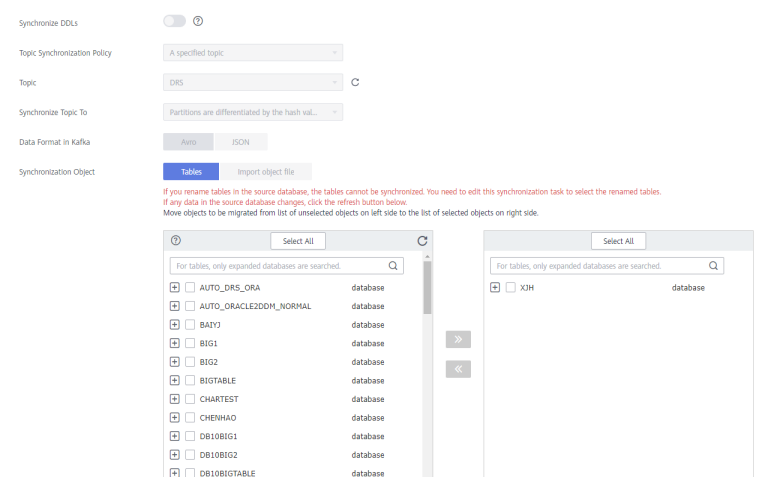



Table 5-29 Synchronization mode and object

Parameter	Description
Synchronize DDLs	Whether to synchronize DDLs. If Synchronize DDLs is enabled and Partitions are identified by the hash values of the primary key is selected, a hash operation is performed on DDLs based on the table name.
Topic Synchronization Policy	Topic synchronization policy. The options are as follows: <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	Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy 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>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.</p> <p>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>
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 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

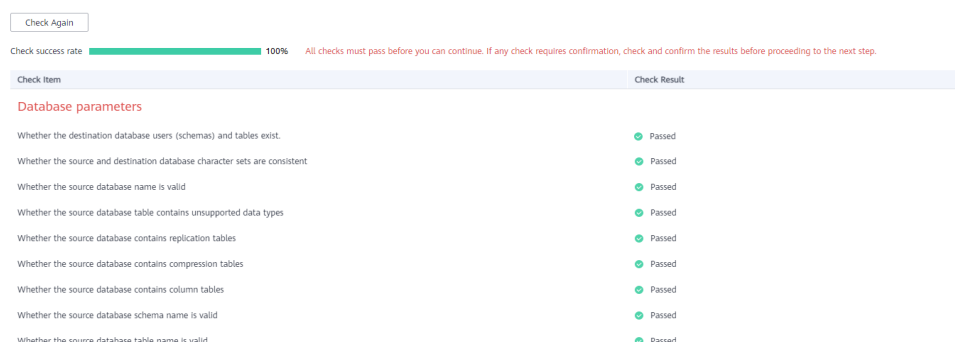
Step 5 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-32 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-33 Task startup settings

The screenshot shows the 'Task startup settings' interface. At the top, there are two tabs: 'Start upon task creation' (selected) and 'Start at a specified time'. Below the tabs, there are several settings:


- Send Notifications:** A toggle switch is turned on. A red note says: 'Please handle exceptions within 48 hours of receiving SMS messages or emails.'
- SMN Topic:** A dropdown menu with a refresh icon and a help icon.
- Synchronization Delay Threshold(s):** A toggle switch is turned off.
- Stop Abnormal Tasks After:** An input field containing the number '14'. A red note says: 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 5-30 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.4 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-31 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	GaussDB primary/standby

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 5-32 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+incremental synchronization and incremental synchronization: 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● Destination database:

Type	Restrictions
	<ul style="list-style-type: none"> <li data-bbox="627 297 1374 427">- 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. <li data-bbox="627 443 1406 573">- 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. <li data-bbox="627 589 1422 741">- 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, 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, xml, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: bfile, xml, sdo_geometry, urowid, interval, and user-defined types. ● The maximum precision supported by timestamp and interval day to second is 6. ● 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. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, and function indexes. ● 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. ● 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. ● 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 the table contains only LOB columns, data inconsistency 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.

Type	Restrictions
	<ul style="list-style-type: none"> • 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 expressions of functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • 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 due to time zone conversion.
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 following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.

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: <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.

Type	Restrictions
Precautions	<ul style="list-style-type: none">● 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.● 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. .● 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.● 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.● If the PDB database is used for synchronization, all PDBs must be enabled during 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.● 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.● 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

Type	Restrictions
	<p data-bbox="624 297 1369 461">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.</p> <ul data-bbox="587 477 1428 1995" style="list-style-type: none"><li data-bbox="587 477 1428 808">● 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.<li data-bbox="587 824 1428 920">● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.<li data-bbox="587 936 1428 1032">● During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored.<li data-bbox="587 1048 1428 1144">● During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails.<li data-bbox="587 1160 1428 1218">● Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent.<li data-bbox="587 1234 1428 1496">● 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.<li data-bbox="587 1512 1428 1608">● 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.<li data-bbox="587 1624 1428 1921">● 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.<li data-bbox="587 1937 1428 1995">● During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During

Type	Restrictions
	<p>incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database.</p> <ul style="list-style-type: none"> ● When a table structure is fully synchronized, only default value constraints of the character string or number type are supported. Default value constraints of the function and sequence types are not supported. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, the following DDL operations are supported: <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. - Database-level synchronization supports create table. (Table definitions cannot contain functions.) ● 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. ● Full and incremental synchronization does not support the STRING extended data type (EXTENDED DATA TYPE) of Oracle 12c or later. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the source database uses XStream for log reading, you need to insert data during incremental synchronization. Otherwise, the incremental synchronization delay keeps increasing. ● When DRS reads incremental Oracle logs in XStream mode, an Oracle XStream outbound is created on the source database. If another XStream outbound that reports an error already exists on the source database, the new XStream outbound created by DRS may be affected. As a result, data cannot be read. You can run the following command in the source database to check whether there are other XStream outbound errors: <pre style="background-color: #f0f0f0; padding: 5px;">select o.server_name, c.capture_name, c.error_message from SYS.ALL_XSTREAM_OUTBOUND O, SYS.ALL_CAPTURE C where o.capture_name = c.capture_name</pre> ● When editing the task to add a new table, ensure that transactions of the new table have been committed.

Type	Restrictions
	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.

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 5-34 Synchronization task information

Table 5-33 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-35 Synchronization instance details

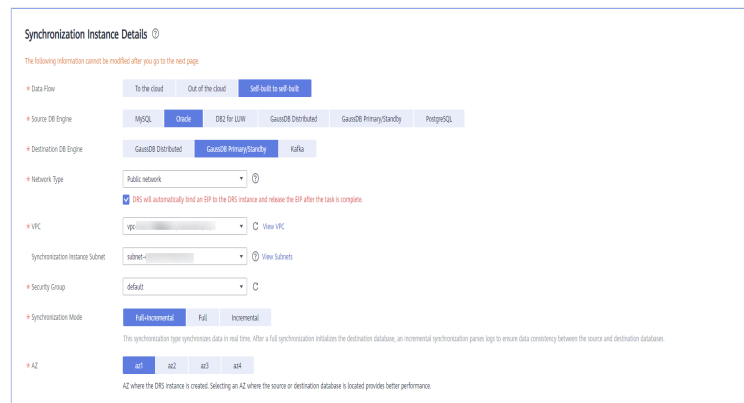


Table 5-34 Synchronization instance settings

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
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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	Available options: Full , Incremental , and Full +Incremental .
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Projects and Tags

Figure 5-36 Enterprise projects and tags




Table 5-35 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 5-37 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
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 5-36 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 not enabled, your data may be at risk.

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 🔒

This button is available only after t

Table 5-37 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 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.</p>
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 5-39 Synchronization mode


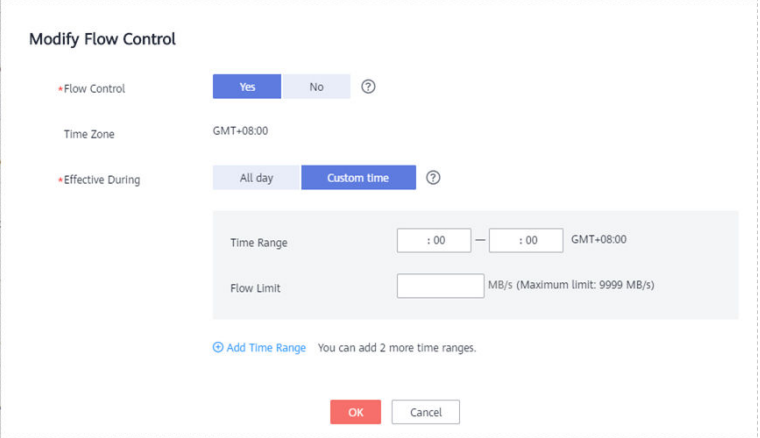



Table 5-38 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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	<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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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, view the parameters you set for full synchronization or incremental synchronization in [Step 2](#). The default parameter values cannot be modified. Click **Next**.


Figure 5-41 Full synchronization

Full Synchronization Settings


* Synchronization Object Type Table structure Data Constraint

* Stream Mode 

* Concurrent Export Tasks 

* Concurrent Import Tasks 

* Import Mode COPY INSERT
 The copy mode improves write efficiency.

* Rows per Shard 


Synchronization Point 

Table 5-39 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Available options: Table structure , Data , and Constraint Data is mandatory.	All options
Stream Mode	In this mode, data is written to the destination database by shard. Data is submitted once for each shard. Otherwise, The data is submitted once every 16 MB of data is written to the destination database.	Enabled
Concurrent Export Tasks	The number of threads that can be started to export data from the Oracle database.	8
Concurrent Import Tasks	The number of threads that can be started to concurrently write data to GaussDB.	8
Import Mode	<ul style="list-style-type: none">• COPY With the COPY interface, writing efficiency is relatively high, but task errors may occur due to data conflicts.• INSERT Writing data using the INSERT statement can avoid data conflicts but the writing efficiency is low.	COPY
Rows per Shard	Export tables in shards to improve the efficiency of real-time synchronization. <ul style="list-style-type: none">• 0 All tables are not sharded, and each table is synchronized as a whole.• Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded.• If the table is a partitioned table, each partition is synchronized as a shard instead of being sharded based on the value you set.• If a table has no primary key or unique index, the table is not sharded. In this case, this parameter is meaningless.• If a table is not analyzed in a timely manner, the value of NUM_ROWS in the ALL_TABLES view is empty, and the table is not sharded.	520000

Parameter	Description	Default Value
Synchronization Point	Set this parameter to the current system change number (SCN). If you set this parameter to a future SCN or the SCN range where deleted archive logs are located, the task will fail.	The default value is the SCN generated when the task is started.

Figure 5-42 Incremental synchronization

Incremental Data Capture Settings

* Concurrent Log Capture Tasks ?

Capture Start Point

If this parameter is not specified, the start point of the current database is used. (In Oracle database, SCN is used.)

Incremental Replay Settings

* Concurrent Replay Tasks ?

* Replay Policy Automated Manual
After the task is started, the replay component is started automatically.

* Conflict Policy Overwrite Report error Ignore
The destination database data is overwritten by the new data.

Replay Start Point

If this parameter is not specified, the replay start point is the same as the capture start point by default.

Table 5-40 Incremental synchronization settings


Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads for reading Oracle logs. The value ranges from 1 to 16.	1
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.	If this parameter is left blank, the current SCN of the database is used as the start point by default.

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
Replay Policy	<ul style="list-style-type: none"> • Automated After the task is started, the replay component is started automatically. • Manual After a task is started, the replay component needs to be started manually. 	Auto
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
Replay Start Point	Specifies the SCN where the replay starts.	If this parameter is not specified, the replay start point is the same as the capture start point by default.

Step 6 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 5-43 Processing data



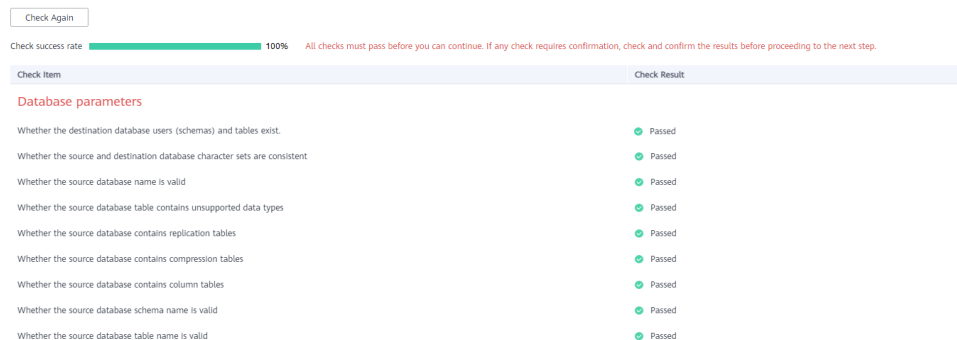
Step 7 On the **Check Task** page, check the synchronization task.

- 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-44 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-45 Task startup settings

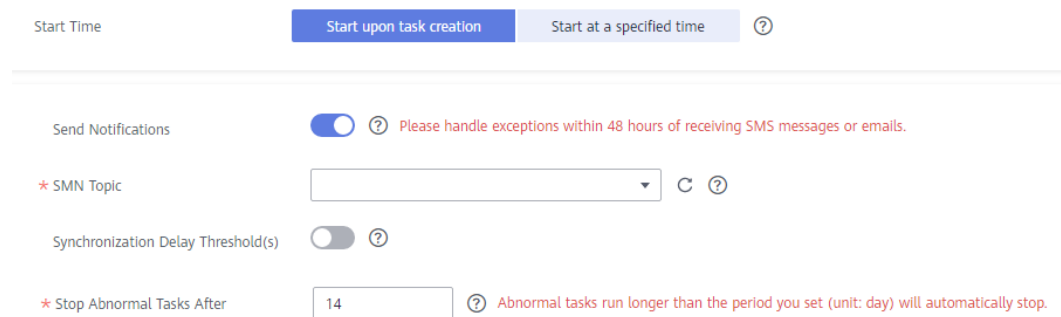



Table 5-41 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.5 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 5-42 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • On-premises databases • ECS databases 	GaussDB distributed

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](#).

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.
 - Start your synchronization task during off-peak hours. 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 migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [What Is the Impact of DRS on 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 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.

Table 5-43 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+incremental synchronization and incremental synchronization: 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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. ● Destination database:

Type	Restrictions
	<ul style="list-style-type: none"><li data-bbox="624 300 1374 427">- 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.<li data-bbox="624 443 1406 571">- 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.<li data-bbox="624 586 1422 741">- 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.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary keys, 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, xmltype, sdo_geometry, urowid, and user-defined types. Incremental synchronization does not support the following column types: XMLTYPE, bfile, xmltype, sdo_geometry, urowid, interval year to month, interval day to second, and user-defined types. ● 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 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. ● Full synchronization and structure migration do not support bitmap indexes, inverted indexes, and function indexes. ● The maximum precision supported by timestamp and interval day to second is 6. ● 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. ● 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 the table contains only LOB columns, data inconsistency 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

Type	Restrictions
	<p>the processing of the empty string in the destination database.</p> <ul style="list-style-type: none"> ● 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 expressions of functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● 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 due to time zone conversion.
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 following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.

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:<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.

Type	Restrictions
Precautions	<ul style="list-style-type: none"> ● 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. ● 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. . ● 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. ● 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. ● If the PDB database is used for synchronization, all PDBs must be enabled during 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. ● 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. ● 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

Type	Restrictions
	<p>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.</p> <ul style="list-style-type: none"> ● 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. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● 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. ● 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 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. ● 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, if the destination database has a constraint with the same name as the source database, the table fails to be created. ● During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During

Type	Restrictions
	<p>incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database.</p> <ul style="list-style-type: none"> ● When a table structure is fully synchronized, only default value constraints of the character string or number type are supported. Default value constraints of the function and sequence types are not supported. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, the following DDL operations are supported: <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. - Database-level synchronization supports create table. (Table definitions cannot contain functions.) ● 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. ● Full and incremental synchronization does not support the STRING extended data type (EXTENDED DATA TYPE) of Oracle 12c or later. ● Full and incremental synchronization does not support hidden columns (UNUSED, INVISIBLE). ● If the source database uses XStream for log reading, you need to insert data during incremental synchronization. Otherwise, the incremental synchronization delay keeps increasing. ● When DRS reads incremental Oracle logs in XStream mode, an Oracle XStream outbound is created on the source database. If another XStream outbound that reports an error already exists on the source database, the new XStream outbound created by DRS may be affected. As a result, data cannot be read. You can run the following command in the source database to check whether there are other XStream outbound errors: <pre style="background-color: #f0f0f0; padding: 5px;">select o.server_name, c.capture_name, c.error_message from SYS.ALL_XSTREAM_OUTBOUND O, SYS.ALL_CAPTURE C where o.capture_name = c.capture_name</pre> ● If the destination database is a GaussDB distributed database, the update operation performed on the source distribution

Type	Restrictions
	<p>column will fail to be synchronized to the GaussDB database during incremental synchronization, causing data inconsistency. Therefore, avoid updating the distribution column.</p> <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.

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 5-46 Synchronization task information

The screenshot shows a web form for creating a synchronization instance. At the top, there is an orange warning banner with the text: "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." Below the banner, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "* Task Name" (a text input field containing "DRS-0131"), and "Description" (a text area with a character count of "0/256"). A small note below the Region field states: "Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region."

Table 5-44 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-47 Synchronization instance details

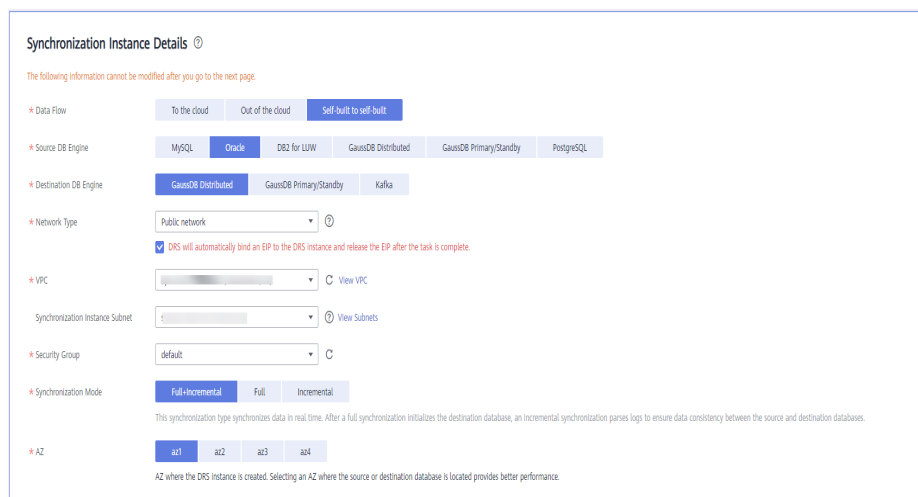


Table 5-45 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
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
VPC	Select an available VPC.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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	Available options: Full , Incremental , and Full +Incremental .
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Projects and Tags

Figure 5-48 Enterprise projects and tags

Table 5-46 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> – This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. – 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 ends.

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 5-49 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 5-47 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 not enabled, your data may be at risk.

Figure 5-50 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...

Table 5-48 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 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.</p>
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 5-51 Synchronization mode


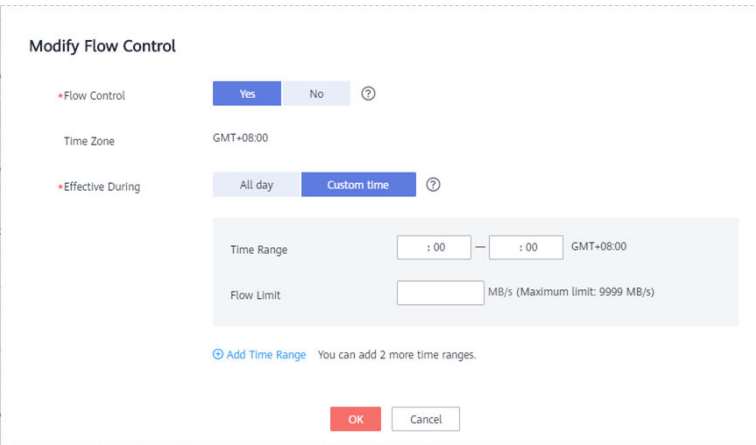



Table 5-49 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 5-52 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. 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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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, view the parameters you set for full synchronization or incremental synchronization in [Step 2](#). The default parameter values cannot be modified. Click **Next**.


Figure 5-53 Full synchronization

Full Synchronization Settings


* Synchronization Object Type Table structure Data Constraint

* Stream Mode 

* Concurrent Export Tasks 

* Concurrent Import Tasks 

* Import Mode COPY INSERT
 The copy mode improves write efficiency.

* Rows per Shard 


Synchronization Point 

Table 5-50 Full synchronization settings

Parameter	Description	Default Value
Synchronization Object Type	Available options: Table structure , Data , and Constraint Data is mandatory.	All options
Stream Mode	In this mode, data is written to the destination database by shard. Data is submitted once for each shard. Otherwise, The data is submitted once every 16 MB of data is written to the destination database.	Enabled
Concurrent Export Tasks	The number of threads that can be started to export data from the Oracle database.	8
Concurrent Import Tasks	The number of threads that can be started to concurrently write data to GaussDB.	8
Import Mode	<ul style="list-style-type: none"> • COPY With the COPY interface, writing efficiency is relatively high, but task errors may occur due to data conflicts. • INSERT Writing data using the INSERT statement can avoid data conflicts but the writing efficiency is low. 	COPY
Rows per Shard	<p>Export tables in shards to improve the efficiency of real-time synchronization.</p> <ul style="list-style-type: none"> • 0 All tables are not sharded, and each table is synchronized as a whole. • Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. • If the table is a partitioned table, each partition is synchronized as a shard instead of being sharded based on the value you set. • If a table has no primary key or unique index, the table is not sharded. In this case, this parameter is meaningless. • If a table is not analyzed in a timely manner, the value of NUM_ROWS in the ALL_TABLES view is empty, and the table is not sharded. 	520000

Parameter	Description	Default Value
Synchronization Point	Set this parameter to the current system change number (SCN). If you set this parameter to a future SCN or the SCN range where deleted archive logs are located, the task will fail.	The default value is the SCN generated when the task is started.

Figure 5-54 Incremental synchronization

Incremental Data Capture Settings

* Concurrent Log Capture Tasks ?

Capture Start Point

If this parameter is not specified, the start point of the current database is used. (In Oracle database, SCN is used.)

Incremental Replay Settings

* Concurrent Replay Tasks ?

* Replay Policy Automated Manual
After the task is started, the replay component is started automatically.

* Conflict Policy Overwrite Report error Ignore
The destination database data is overwritten by the new data.

Replay Start Point

If this parameter is not specified, the replay start point is the same as the capture start point by default.

Table 5-51 Incremental synchronization settings


Parameter	Description	Default Value
Concurrent Log Capture Tasks	The number of concurrent threads for reading Oracle logs. The value ranges from 1 to 16.	1
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.	If this parameter is left blank, the current SCN of the database is used as the start point by default.

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
Replay Policy	<ul style="list-style-type: none"> • Automated After the task is started, the replay component is started automatically. • Manual After a task is started, the replay component needs to be started manually. 	Auto
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
Replay Start Point	Specifies the SCN where the replay starts.	If this parameter is not specified, the replay start point is the same as the capture start point by default.

Step 6 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 5-55 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-56 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-57 Task startup settings





Table 5-52 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.6 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 5-53 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 5-54](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-54 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. • 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>CAUTION 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>

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 5-55](#). 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.

Table 5-55 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 [What Is the Impact of DRS on 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 5-56 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. • 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. - 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. • 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.
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.

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 5-58 Synchronization task information

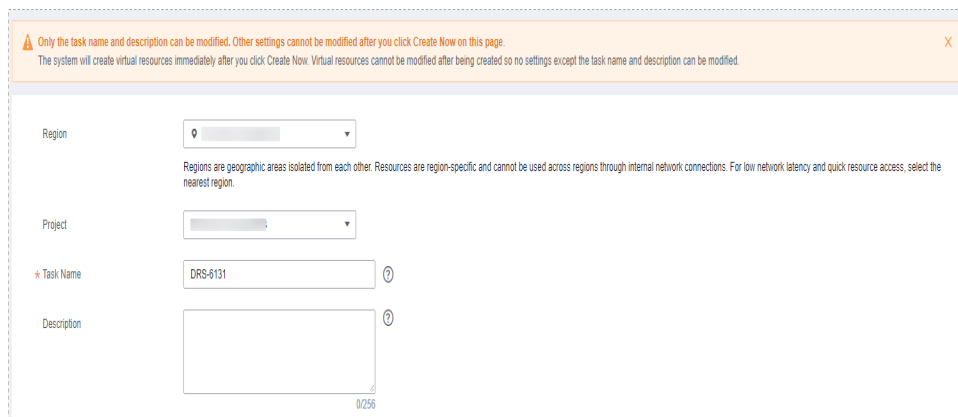


Table 5-57 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-59 Synchronization instance details

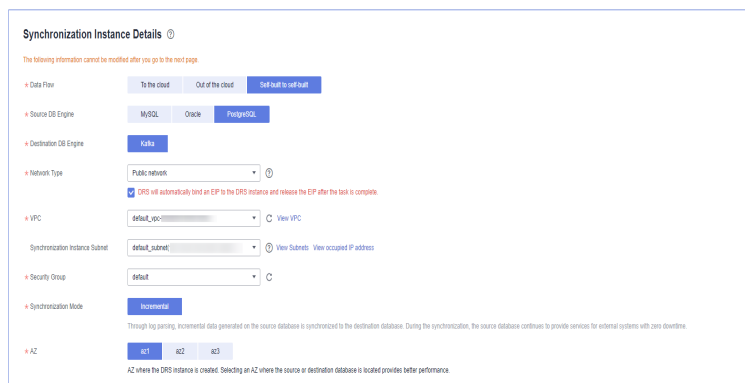


Table 5-58 Synchronization instance settings

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 Subnet 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 ensure that there are available IP addresses. 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	– 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.
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.

- DRS instance specifications

Figure 5-60 Specifications



Table 5-59 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 5-61 Enterprise projects and tags




Table 5-60 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 5-62 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 Name

Database Username

Database Password

SSL Connection

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

Table 5-61 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 5-63 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Connection Method


 Test successful

Table 5-62 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 5-64 Synchronization mode

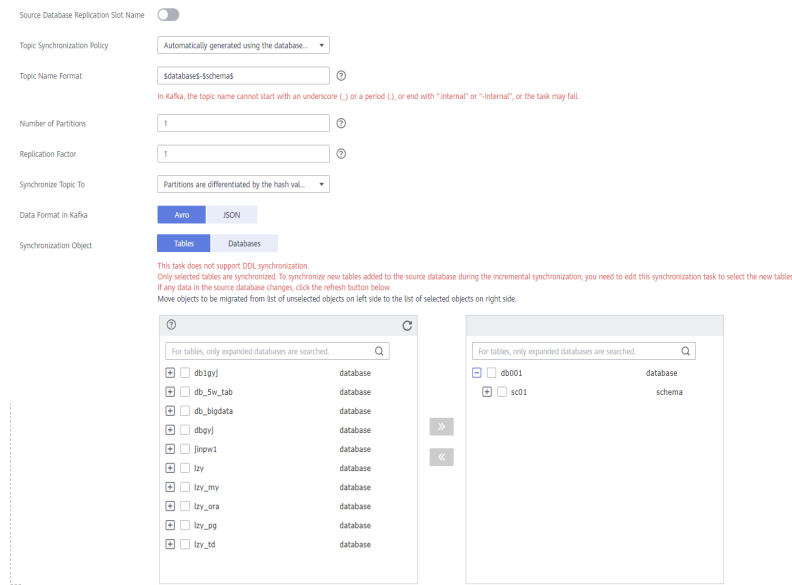



Table 5-63 Synchronization object

Parameter	Description
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.

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"> ● 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>

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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-65 Task 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-66 Task startup settings

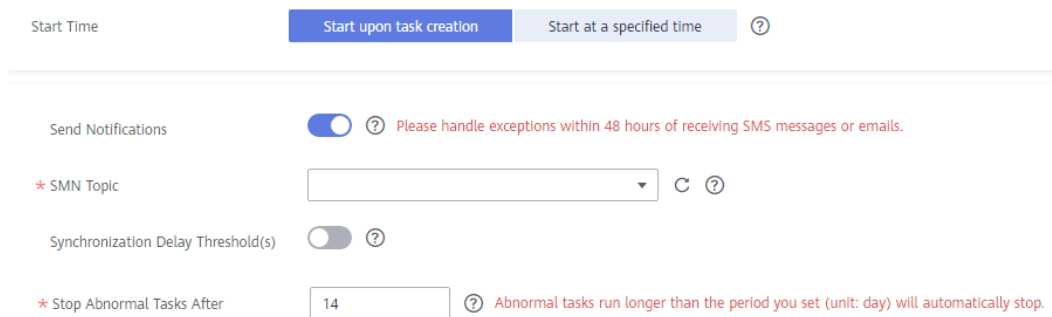



Table 5-64 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.7 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 5-65 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 5-66 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-66 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 5-67](#). 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.

Table 5-67 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>	<p>The user has the sysadmin role or the following minimum permissions:</p> <p>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.</p>
Destination database user	<p>The user must have the DBA role or the following minimum permissions:</p> <p>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:</p> <p>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:</p> <p>CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role</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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 5-68 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 or GB18030 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 synchronization may be inconsistent, 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"> - 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. - 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. - 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.
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.

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. ● 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 <i>GaussDB Developer Guide</i>. ● 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 and unique key 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 an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task 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.

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 5-67 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]

Project: [Dropdown]

* Task Name: DRS-6131

Description: [Text Area] 0/256

Table 5-69 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-68 Synchronization instance details

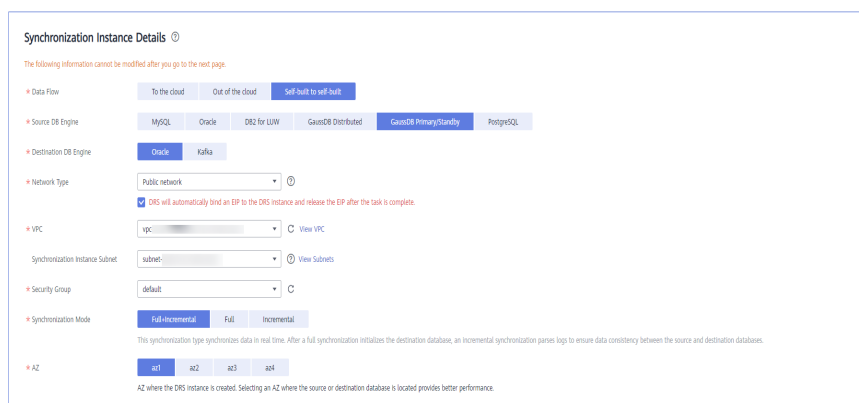


Table 5-70 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
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.

Parameter	Description
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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>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.
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.

- Enterprise Projects and Tags

Figure 5-69 Enterprise projects and tags




Table 5-71 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 5-70 Source database information

Source Database

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

Table 5-72 Source database settings

Parameter	Description
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 5-71 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 5-73 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 not enabled, 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 5-72 Synchronization mode

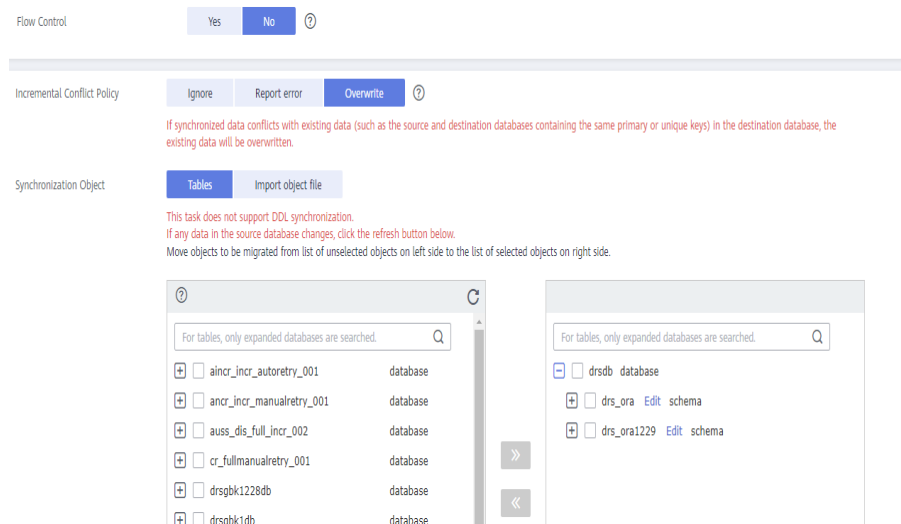
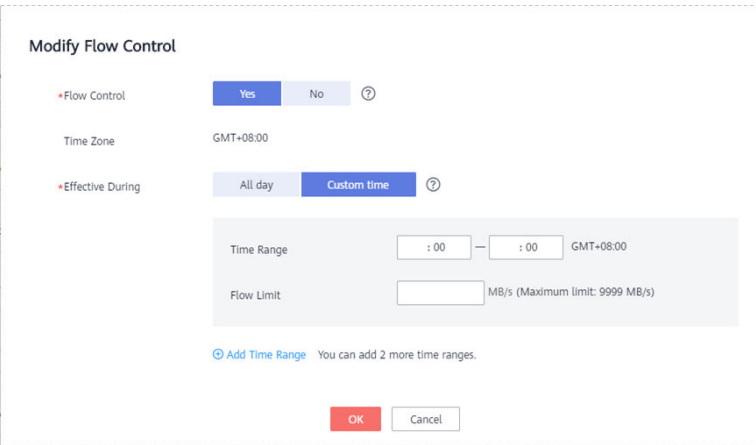



Table 5-74 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 5-73 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 5-74 Parameter settings

Full Synchronization Settings

* Concurrent Export Tasks

* Concurrent Import Tasks

* Rows per Shard

Incremental Replay Settings

* Concurrent Replay Tasks

Table 5-75 Full synchronization settings

Parameter	Description	Default Value
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

Table 5-76 Incremental synchronization settings

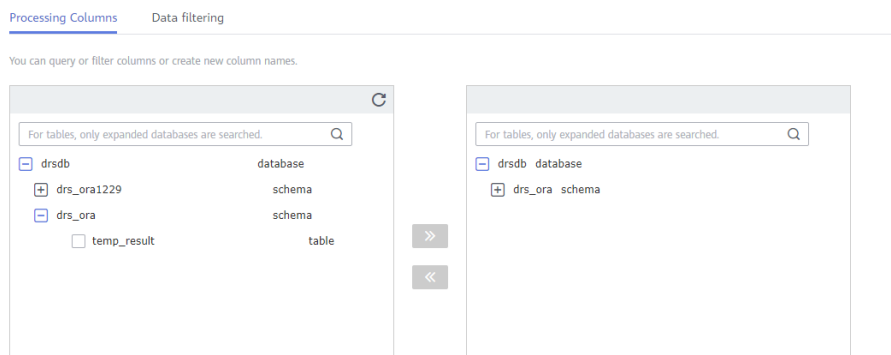
Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

Step 6 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](#).

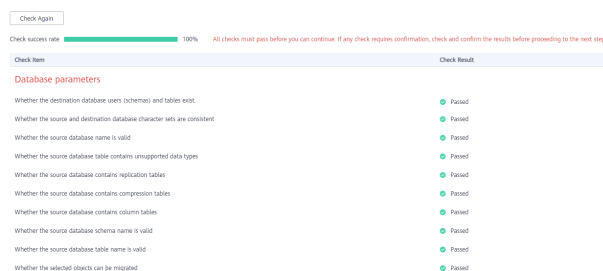
Figure 5-75 Processing data



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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 5-76 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 8 On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-77 Task startup settings





Table 5-77 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.8 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 5-78 Supported databases

Source DB	Destination DB
GaussDB primary/standby	Kafka 0.11 or later

Supported Synchronization Objects

Table 5-79 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-79 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: INTEGER, TINYINT, SMALLINT, BIGINT, 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"> - 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - 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 5-80](#). 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.

Table 5-80 Database user permission

Type	Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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.

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 [What Is the Impact of DRS on 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 5-81 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. ● Other notes: <ul style="list-style-type: none"> - The source database must be a GaussDB primary/standby 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. - 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.

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. • During synchronization of table-level objects, renaming tables is not recommended. • 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.
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> – When an 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 real-time incremental synchronization task, you may need to manually delete residual replication slots from the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby.

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](#).

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 5-78 Synchronization task information

Table 5-82 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 !=<>'&"\

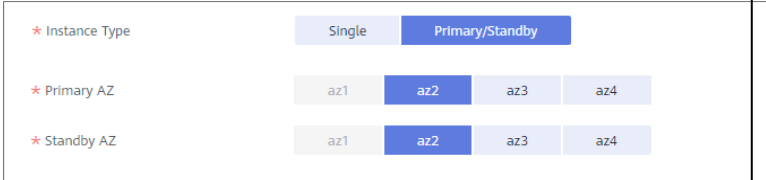
- Synchronization instance details

Figure 5-79 Synchronization instance details

Table 5-83 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 Kafka .
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby. Note that if Primary/Standby is selected, data in tables without primary keys may be duplicate or inconsistent when a primary/standby switchover is performed.</p> <ul style="list-style-type: none"> - Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. - Single: The single-node deployment architecture is used, which is cost-effective.
VPC	Select an available VPC.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet 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 ensure that there are available IP addresses. 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	<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
AZ	<p>Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-80 AZ</p> 

- DRS instance specifications

Figure 5-81 Specifications



Table 5-84 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 5-82 Enterprise projects and tags

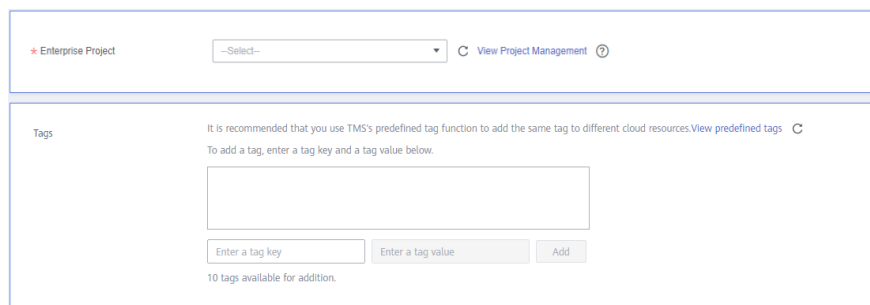


Table 5-85 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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">- This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- 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 ends.

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 5-83 Source database information

Source Database

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

Table 5-86 Source database settings

Parameter	Description
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 5-84 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Connection Method

✔ Test successful

Table 5-87 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 5-85 Synchronization mode

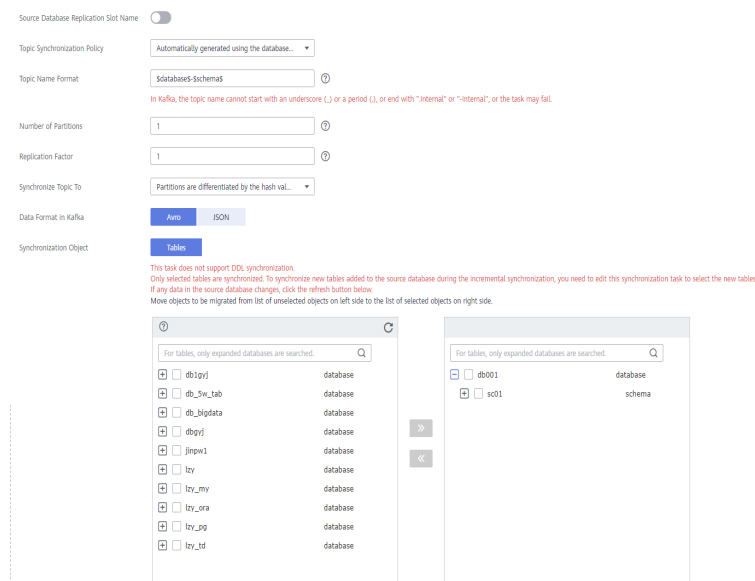



Table 5-88 Synchronization Object

Parameter	Description
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 .

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	<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-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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-86 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-87 Task startup settings

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

Send Notifications ? Please handle exceptions within 48 hours of receiving SMS messages or emails.

* SMN Topic ?

Synchronization Delay Threshold(s) ?


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

Table 5-89 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.9 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 5-90 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 5-91 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-91 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, 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - 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 5-92](#). 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 5-92 Database user permission

Type	Full	Incremental	Full+Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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 has the sysadmin role or the following minimum permissions: 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.

Type	Full	Incremental	Full+Incremental
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

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.
 - Start your synchronization task during off-peak hours. 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 [What Is the Impact of DRS on 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 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 5-93 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 or GB18030 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 synchronization may be inconsistent, 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"> - 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. - During real-time synchronization, the consistency of distributed transactions is not ensured. - 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. - 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.
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.

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. ● 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 <i>GaussDB Developer Guide</i>. ● 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 an incremental or full+incremental synchronization is complete, the streaming replication slot created by the task 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.

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 5-88 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. X

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 ⓘ

Description: [Text area] ⓘ

0/256

Table 5-94 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-89 Synchronization instance details

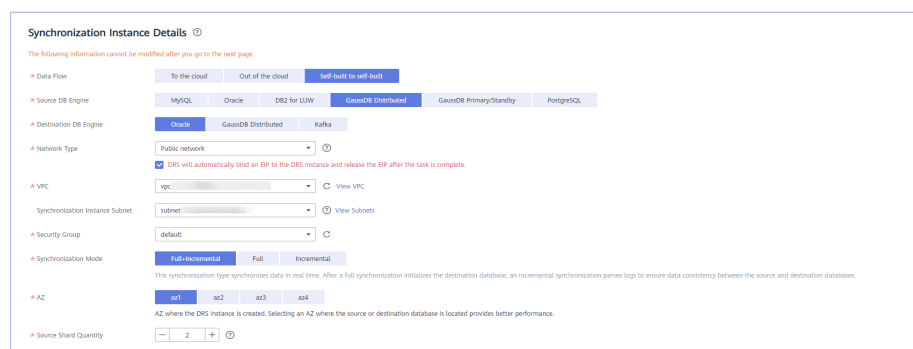


Table 5-95 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Oracle .
Network Type	Public network is used as an example. Available options: Public network and VPN or Direct Connect
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 Subnet 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 ensure that there are available IP addresses. 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>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.
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.
Source Shard Quantity	<p>The value must be the same as the number of DNs in the distributed source database. Run the following SQL statement to check the value of Source Shard Quantity:</p> <pre>SELECT count(DISTINCT node_name) FROM pgxc_node WHERE node_type = 'D';</pre>

- Enterprise Projects and Tags

Figure 5-90 Enterprise projects and tags




Table 5-96 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 5-91 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

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

Table 5-97 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	<p>You can log in to the CN and run the following SQL statement to query the DN IP address:</p> <pre>SELECT node_name as dn_name, string_agg(node_host ':' node_port, ',' order by node_type) AS dn_ip FROM pgxc_node WHERE node_type != 'C' GROUP BY node_name;</pre> <p>The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.</p>

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 5-92 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 5-98 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 not enabled, 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 5-93 Synchronization mode

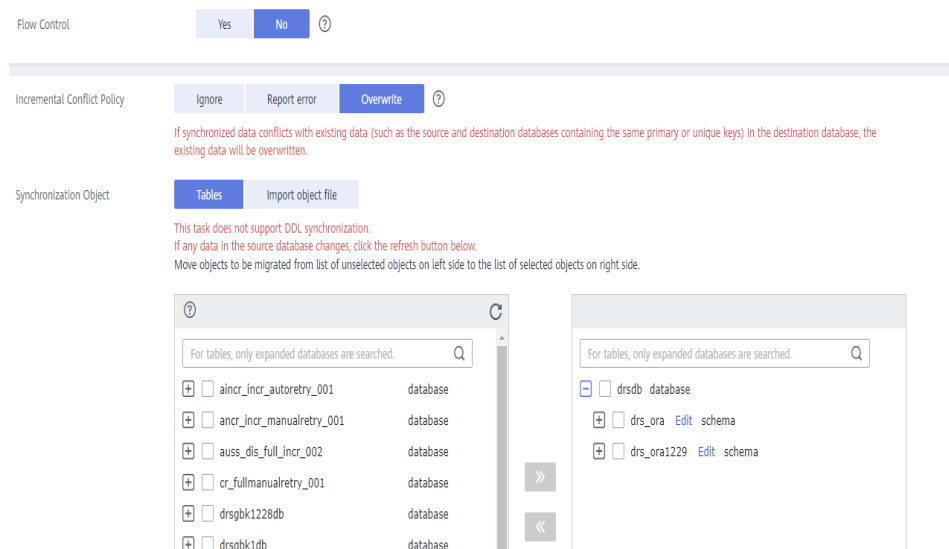
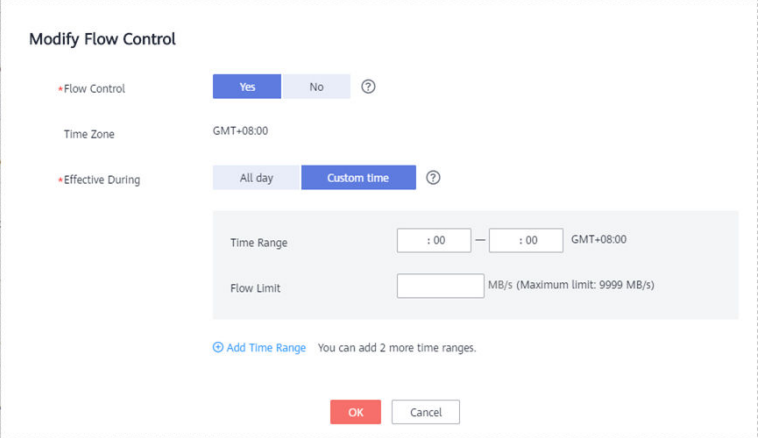



Table 5-99 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 5-94 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.
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 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 full+incremental synchronization selected in [Step 2](#) and click **Next**.

Figure 5-95 Parameter settings

Full Synchronization Settings

* Concurrent Export Tasks ?

* Concurrent Import Tasks ?

* Rows per Shard ?

Incremental Replay Settings

* Concurrent Replay Tasks ?

Table 5-100 Full synchronization settings

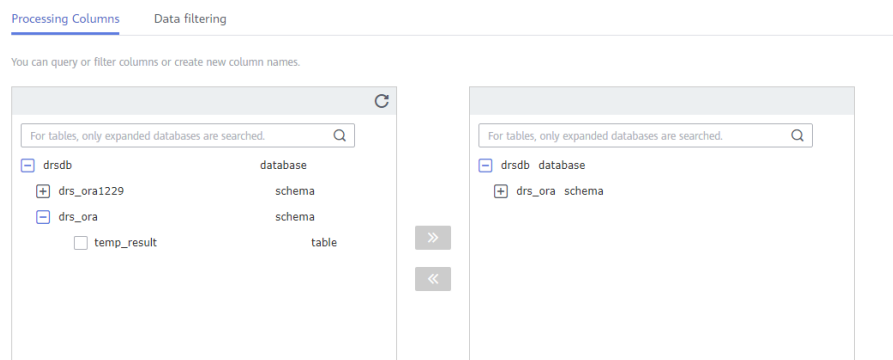
Parameter	Description	Default Value
Concurrent Export Tasks	Number of export threads. Value range: 1 to 16. A larger value indicates higher load on the source database.	8
Concurrent Import Tasks	Number of import threads. Value range: 1 to 16. A larger value indicates higher load on the destination database.	8
Rows per Shard	Value range: 0 or 520,000 to 1,000,000,000 <ul style="list-style-type: none"> 0 All tables are not sharded, and each table is synchronized as a whole. Other values The table is sharded based on the specified value (or the primary key column). If the number of records in a table is less than the value of this parameter, the table is not sharded. 	520000

Table 5-101 Incremental synchronization settings

Parameter	Description	Default Value
Concurrent Replay Tasks	Number of concurrent threads for data replay. Value range: 1 to 64. Incremental data is concurrently written to the destination database. A larger value indicates higher load on the destination database.	64

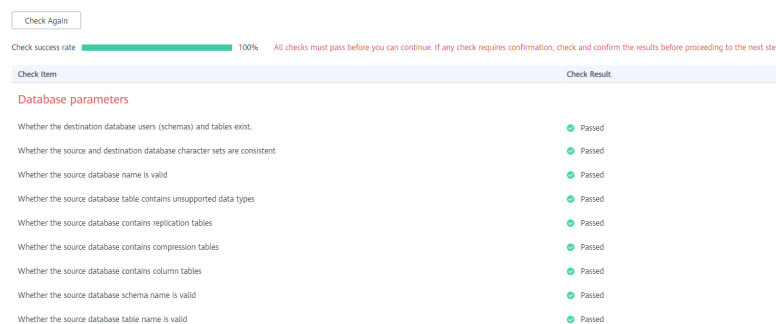
- Step 6** On the **Process Data** page, set the filtering rules for data processing.
- If data processing is not required, click **Next**.
 - If you need process columns or filter out data, set the corresponding rules by referring to **Processing Data**.

Figure 5-96 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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 5-97 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 8** On the displayed page, specify **Start Time**, **Send Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-98 Task startup settings





Table 5-102 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

5.10 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 5-103 Supported databases

Source DB	Destination DB
GaussDB distributed	Kafka 0.11 or later

Supported Synchronization Objects

Table 5-104 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-104 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: INTEGER, TINYINT, SMALLINT, BIGINT, 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"> - 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, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - 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 5-105](#). 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 5-105 Database user permission

Type	Incremental
Source database user	The user has the sysadmin role or the following minimum permissions: 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.

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 [What Is the Impact of DRS on 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 5-106 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. ● 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. - 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 change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • During synchronization of table-level objects, renaming tables is not recommended. • 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 <i>GaussDB Developer Guide</i>.
Stopping a task	<ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> - When an 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 real-time incremental synchronization task, you may need to manually delete residual replication slots from the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed.

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 5-99 Synchronization task information

Table 5-107 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 5-100 Synchronization instance details

Table 5-108 Synchronization instance settings

Parameter	Description
Data Flow	Choose Self-built to self-built .

Parameter	Description
Source DB Engine	Select GaussDB Distributed .
Destination DB Engine	Select Kafka .
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 Subnet 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 ensure that there are available IP addresses. 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	– 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.
AZ	Select the AZ where you want to create the DRS instance. Selecting the one housing the source or destination database can provide better performance.
Source Shard Quantity	The value must be the same as the number of DNs in the distributed source database. Run the following SQL statement to check the value of Source Shard Quantity : <code>SELECT count(DISTINCT node_name) FROM pgxc_node WHERE node_type = 'D';</code>

- DRS instance specifications

Figure 5-101 Specifications



Table 5-109 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 5-102 Enterprise projects and tags




Table 5-110 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. 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 ends.

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 5-103 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

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

Table 5-111 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.


Parameter	Description
DN IP Address or Domain Name	<p>You can log in to the CN and run the following SQL statement to query the DN IP address:</p> <pre>SELECT node_name as dn_name, string_agg(node_host ':' node_port, ',' order by node_type) AS dn_ip FROM pgxc_node WHERE node_type != 'C' GROUP BY node_name;</pre> <p>The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel.</p>

 **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 5-104 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Connection Method


 Test successful

Table 5-112 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 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.</p>
Connection Method	<p>Available options: PLAINTEXT, SSL, SASL_PLAINTEXT, and SASL_SSL. For details, see Kafka Authentication.</p>

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

Figure 5-105 Synchronization mode

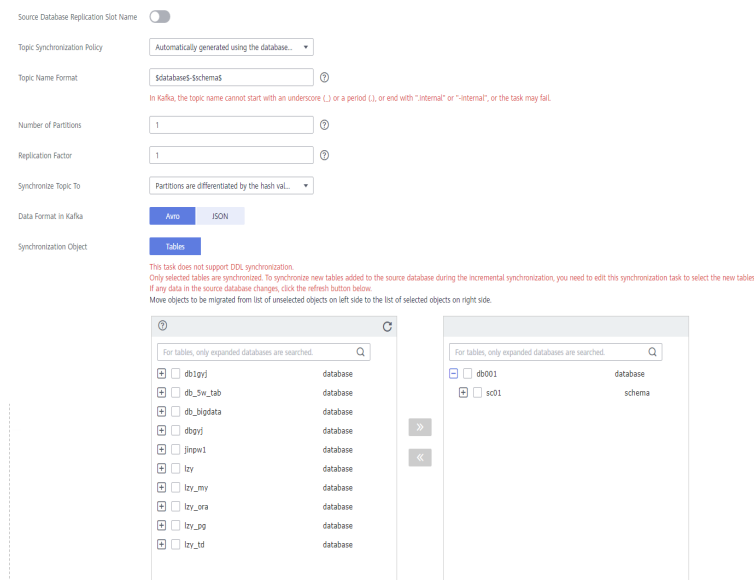



Table 5-113 Synchronization object

Parameter	Description
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"> • 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>

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>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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 5-106 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 5-107 Task startup settings

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

Send Notifications ? Please handle exceptions within 48 hours of receiving SMS messages or emails.

* SMN Topic ? ↻ ?

Synchronization Delay Threshold(s) ?


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

Table 5-114 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.
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 .
Synchronization Delay Threshold	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 Notification.

Parameter	Description
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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

6 Task Management

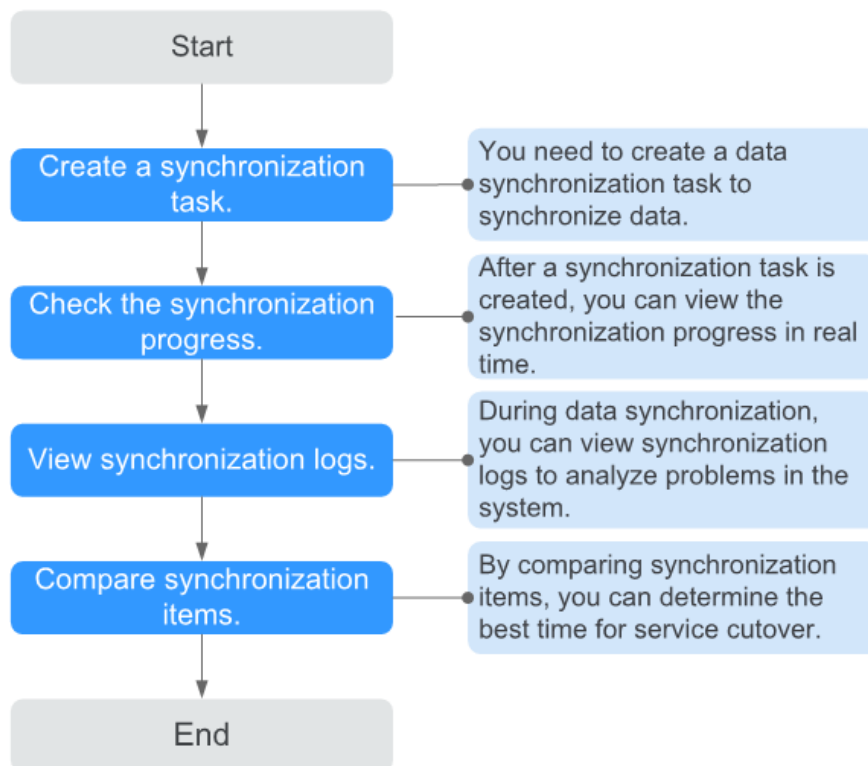
6.1 Creating a Synchronization Task

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 6-1 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.

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](#).

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 6-2 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: [Region] ▾
Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Project: [Project] ▾

* Task Name: DRS-6131 ⓘ

Description: [Description] ⓘ
0/256

Table 6-1 Task and recipient description

Parameter	Description
Region	The region where the synchronization 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 6-3 Synchronization instance information

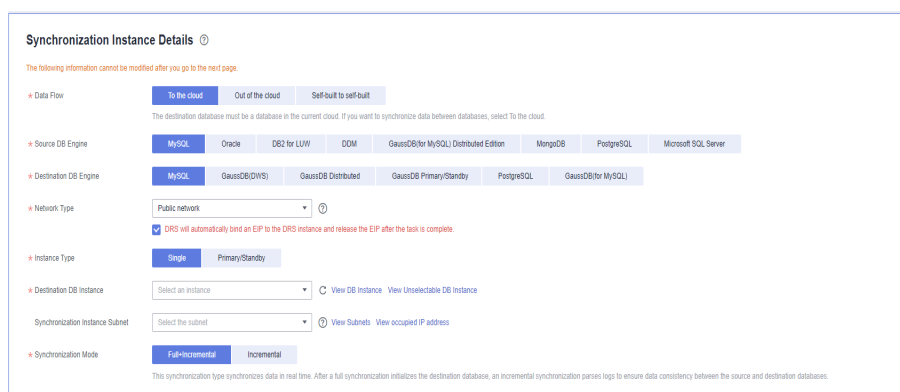


Table 6-2 Synchronization instance settings

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	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect

Parameter	Description
Instance Type	<p>DRS instance type. The value can be Single or Primary/Standby.</p> <ul style="list-style-type: none"> - Primary/Standby: This architecture provides HA, improving the reliability of DRS instances. After a primary/standby instance is created, DRS creates two subtasks, each running on the primary and standby nodes. If the subtask on the primary node fails, DRS automatically starts the subtask on the standby node to continue the synchronization. - Single: The single-node deployment architecture is used, which is cost-effective. <p>This option is available only in specific scenarios. For details, see Performing a Primary/Standby Switchover.</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 Subnet 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 ensure that there are available IP addresses. 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. <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> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

Parameter	Description
AZ	<p>This parameter is available only when you select primary/standby for Instance Type. It indicates the AZ where the DRS instance is created. If the source or destination database is in the same AZ as the DRS instance, you can get better performance.</p> <p>If Instance Type is set to primary/standby, you can specify Primary AZ and Standby AZ.</p> <p>Figure 6-4 AZ</p>

- DRS instance specifications

Figure 6-5 Specifications



Table 6-3 Specifications

Parameter	Description
Specifications	<p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE Currently, DRS supports specification upgrade only in MySQL-to-MySQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded. For details, see Changing Specifications.</p>

- Enterprise Projects and Tags

Figure 6-6 Enterprise projects and tags

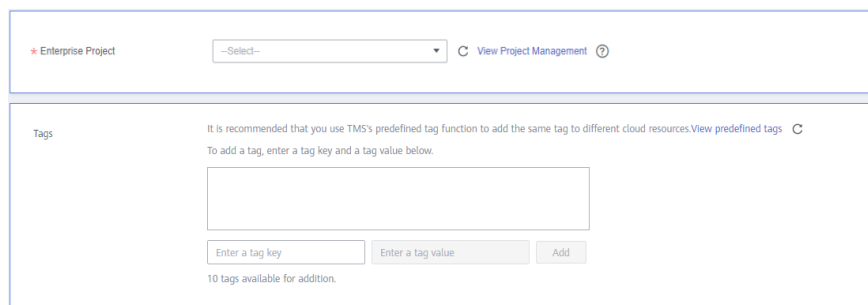


Table 6-4 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default. For more information about enterprise project, see <i>Enterprise Management User Guide</i>.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Management 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"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - 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 ends.

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 6-7 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.

IP Address or Domain Name

Port

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 6-5 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.
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 not enabled, 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 6-8 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection


 Test successful

Table 6-6 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 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	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 not enabled, 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 6-9 Synchronization mode

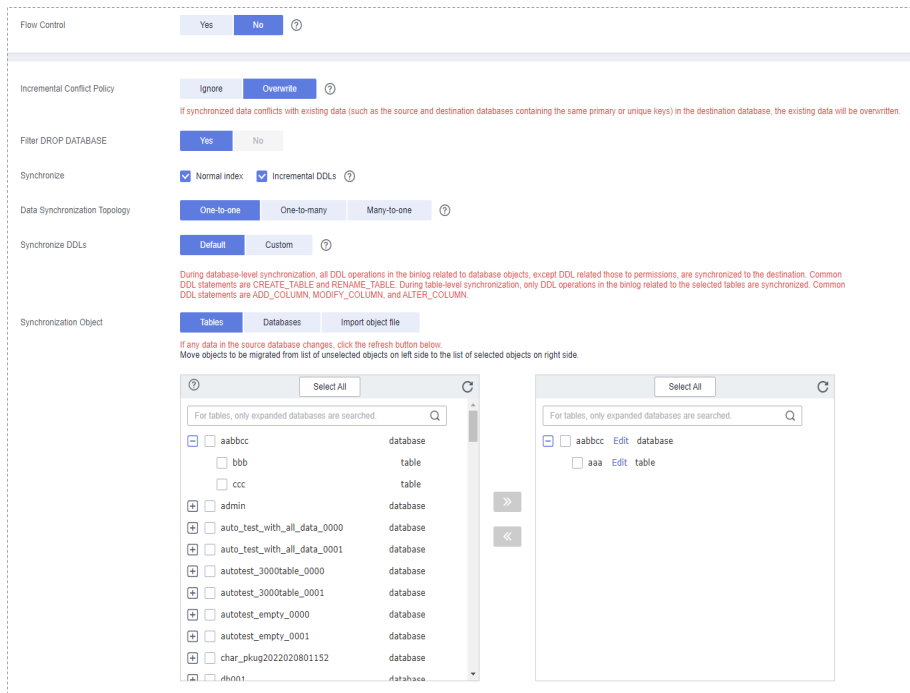
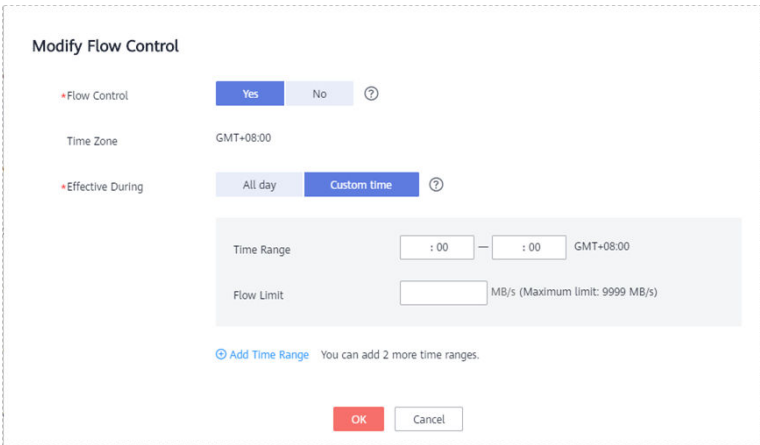



Table 6-7 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none"> Yes You can customize the maximum migration speed. 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 All day. A maximum of three 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 6-10 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. • 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 data based on service requirements.</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 start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</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> <p>NOTE Data Synchronization Topology can be selected only for whitelisted users. 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.</p>

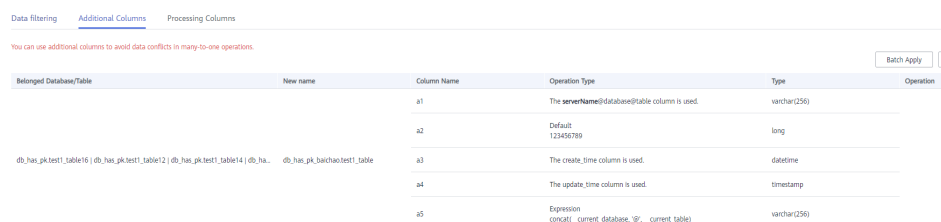
Parameter	Description
Synchronize DDLs	<p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default <ul style="list-style-type: 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. <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 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. 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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and 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 Column**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 6-11 Processing data

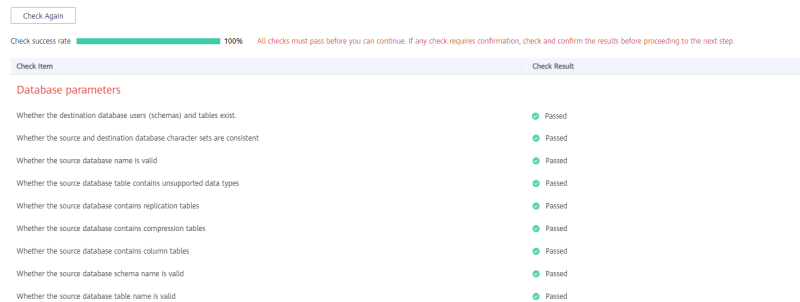


Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
		a1	The serverName@database@table column is used.	varchar(256)	
		a2	Default 123456789	long	
db_has_pk_test1_table16 db_has_pk_test1_table12 db_has_pk_test1_table14 db_has_pk_test1_table18		a3	The create_time column is used.	datetime	
		a4	The update_time column is used.	timestamp	
		a5	Expression concat(_current_database, '...', _current_table)	varchar(256)	

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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 6-12 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, 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 6-13 Task startup settings

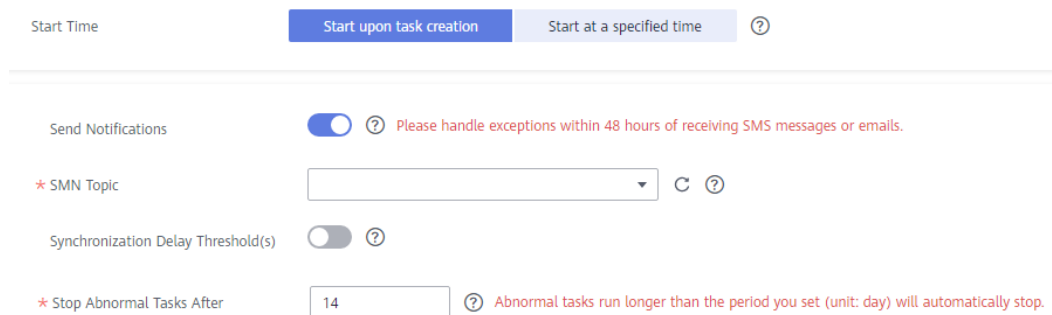



Table 6-8 Task startup settings

Parameter	Description
Started 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	SMN topic. This parameter is optional. If an exception occurs during synchronization, the system will send a notification to the specified recipients.

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 .
Synchronization Delay Threshold	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 Notification.
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 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, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

Helpful Links

- [Supported Databases](#)
- [Preparations](#)
- [Synchronization Overview](#)
- [Data Synchronization Topologies](#)

6.2 Querying the Synchronization Progress

This section describes how to check the synchronization progress.

- During a full synchronization, DRS displays the progress overview. You can view the structure, data, and index migration progress. When the progress reaches 100%, the synchronization is complete. The synchronization of data and indexes is relatively slow.
- 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.
 - 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, the data in the source and destination databases is synchronized 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


Frequent DDL operations, ultra-large transactions, and network problems may result in excessive synchronization delay.

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

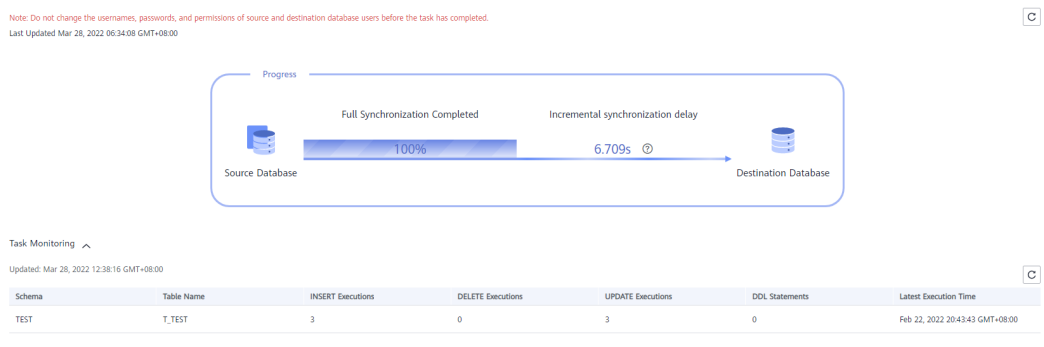
The displayed position information is updated every 10 seconds.

Figure 6-14 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 6-15 Task monitoring

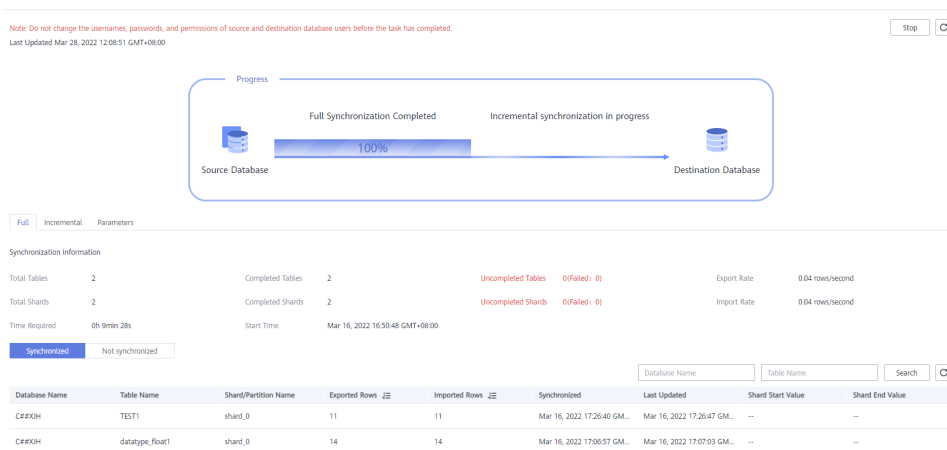


NOTE

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 synchronization scenario, search for the detailed synchronization object information on the **Synchronization Progress** tab. The keywords for searching synchronization objects are case sensitive.

Figure 6-16 Synchronization details



----End

6.3 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.

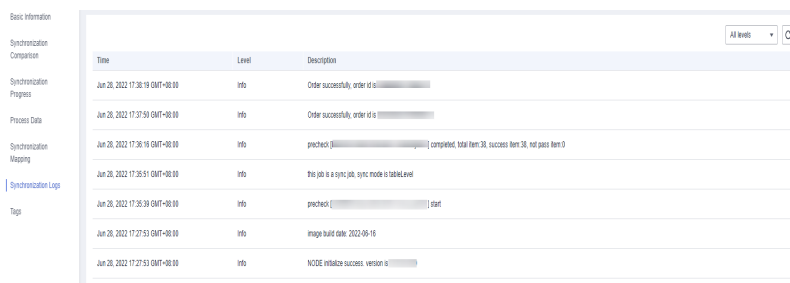
Procedure

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 6-17 Synchronization logs



Time	Level	Description
Jun 28, 2022 17:28:19 GMT+08:00	Info	Order successfully, order id is [redacted]
Jun 28, 2022 17:27:59 GMT+08:00	Info	Order successfully, order id is [redacted]
Jun 28, 2022 17:27:58 GMT+08:00	Info	precheck [redacted] completed, total item:28, success item:28, not pass item:0
Jun 28, 2022 17:26:51 GMT+08:00	Info	This job is a sync-job, sync mode is table-level
Jun 28, 2022 17:26:38 GMT+08:00	Info	precheck [redacted] start
Jun 28, 2022 17:27:53 GMT+08:00	Info	image build date: 2022-09-16
Jun 28, 2022 17:27:53 GMT+08:00	Info	NODE initlize success, version is [redacted]

----End

6.4 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.
- 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 6-9 Supported comparison methods

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
To the cloud	MySQL->MySQL	Yes	Yes	Yes	No
To the cloud	MySQL->PostgreSQL	Yes	Yes	No	No
To the cloud	MySQL -> GaussDB distributed	Yes	Yes	No	No
To the cloud	MySQL -> GaussDB primary/ standby	Yes	Yes	No	No
To the cloud	MySQL->GaussDB(for MySQL)	Yes	Yes	Yes	No
To the cloud	MySQL->GaussDB(DWS)	Yes	Yes	No	No
To the cloud	PostgreSQL->PostgreSQL	Yes	Yes	Yes	Yes
To the cloud	PostgreSQL->GaussDB(DWS)	Yes	Yes	No	No
To the cloud	PostgreSQL -> GaussDB primary/standby	Yes	Yes	Yes	No
To the cloud	PostgreSQL -> GaussDB distributed	Yes	Yes	Yes	No
To the cloud	DDM->MySQL	Yes	Yes	No	No
To the cloud	DDM->GaussDB(DWS)	No	Yes	Yes	No

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
To the cloud	DDM->DDM	Yes	Yes	No	No
To the cloud	Oracle->MySQL	Yes	Yes	Yes	No
To the cloud	Oracle->GaussDB(for MySQL)	Yes	Yes	Yes	No
To the cloud	Oracle -> GaussDB primary/ standby	Yes	Yes	Yes	No
To the cloud	Oracle -> GaussDB distributed	Yes	Yes	Yes	No
To the cloud	Oracle->DDM	Yes	Yes	No	No
To the cloud	Oracle->GaussDB(DWS)	Yes	Yes	No	No
To the cloud	Oracle->PostgreSQL	Yes	Yes	No	No
To the cloud	DB2 for LUW -> GaussDB primary/standby	Yes	Yes	No	No
To the cloud	DB2 for LUW -> GaussDB distributed	Yes	Yes	No	No
To the cloud	DB2 for LUW->GaussDB(DWS)	Yes	Yes	No	No
To the cloud	TiDB->GaussDB(for MySQL)	Yes	Yes	No	No

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
To the cloud	Microsoft SQL Server->GaussDB(DWS)	Yes	Yes	No	No
To the cloud	Microsoft SQL Server -> GaussDB primary/standby	Yes	Yes	No	No
To the cloud	Microsoft SQL Server -> GaussDB distributed	Yes	Yes	No	No
To the cloud	Microsoft SQL Server->Microsoft SQL Server	Yes	Yes	No	No
To the cloud	MongoDB->DDS	Yes	Yes	No	No
From the cloud	MySQL->MySQL	Yes	Yes	Yes	No
From the cloud	MySQL->Kafka	No	No	No	No
From the cloud	MySQL->CSS/ES	Yes	Yes	No	No
From the cloud	DDM->MySQL	Yes	Yes	No	No
From the cloud	DDM->Oracle	Yes	Yes	No	No
From the cloud	DDM->Kafka	No	No	No	No
From the cloud	DDS->MongoDB	Yes	Yes	Yes	No

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
From the cloud	PostgreSQL->Kafka	No	No	No	No
From the cloud	GaussDB primary/standby -> MySQL	Yes	Yes	No	No
From the cloud	GaussDB primary/standby -> Oracle	Yes	Yes	No	No
From the cloud	GaussDB primary/standby -> Kafka	No	No	No	No
From the cloud	GaussDB primary/standby -> GaussDB(DWS)	Yes	Yes	Yes	No
From the cloud	GaussDB primary/standby -> GaussDB distributed	Yes	Yes	No	No
From the cloud	GaussDB primary/standby -> GaussDB primary/standby	Yes	Yes	No	No
From the cloud	GaussDB distributed -> MySQL	Yes	Yes	No	No
From the cloud	GaussDB distributed -> Oracle	Yes	Yes	No	No
From the cloud	GaussDB distributed -> GaussDB(DWS)	Yes	Yes	Yes	No
From the cloud	GaussDB distributed -> Kafka	No	No	No	No
From the cloud	GaussDB distributed -> GaussDB distributed	Yes	Yes	No	No

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
From the cloud	GaussDB distributed->GaussDB primary/standby	Yes	Yes	No	No
From the cloud	GaussDB(for MySQL)->MySQL	Yes	Yes	Yes	No
From the cloud	GaussDB(for MySQL)- >GaussDB(DWS)	Yes	Yes	No	No
From the cloud	GaussDB(for MySQL)->Kafka	No	No	No	No
From the cloud	GaussDB(for MySQL)->CSS/ES	Yes	Yes	No	No
From the cloud	GaussDB(for MySQL)->Oracle	Yes	Yes	No	No
Self- built -> Self- built	Oracle->Kafka	No	No	No	No
Self- built -> Self- built	Oracle -> GaussDB primary/ standby	Yes	Yes	Yes	No
Self- built -> Self- built	Oracle -> GaussDB distributed	Yes	Yes	Yes	No
Self- built -> Self- built	MySQL->Kafka	No	No	No	No

Sync hroni zatio n Direc tion	Data Flow	Object -level Compa rison	Row Compa rison	Value Compa rison	Accoun t-level Compa rison
Self-built -> Self-built	MySQL->CSS/ES	Yes	Yes	No	No
Self-built -> Self-built	PostgreSQL->Kafka	No	No	No	No
Self-built -> Self-built	GaussDB primary/standby -> Oracle	Yes	Yes	No	No
Self-built -> Self-built	GaussDB primary/standby -> Kafka	No	No	No	No
Self-built -> Self-built	GaussDB distributed -> Oracle	Yes	Yes	No	No
Self-built -> Self-built	GaussDB distributed -> Kafka	No	No	No	No

Constraints

- A comparison task can be created only when the task is in the incremental phase. When a full task is complete, DRS automatically creates object-level and row comparison tasks.
- 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.

- 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.
- Some data types do not support value comparison. For details, see [Which Data Types Does Not Support Content Comparison?](#)
- To prevent resources from being occupied for a long time, DRS limits the row comparison duration. If the row comparison duration exceeds the threshold, the row comparison task stops automatically. If the source database is a relational database, the row comparison duration is 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration is 30 minutes.
- 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.
- 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 you want to compare values and the DRS task you create supports value comparison, select a large specification for your DRS instance when creating the DRS task.

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 Click the **Synchronization Comparison** tab.

Step 3 Compare synchronization items.


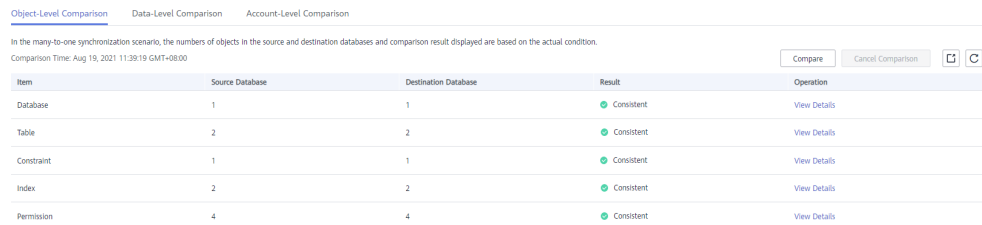
- On the **Object-Level Comparison** tab, click **Compare**. Wait for a while and click  to check whether the comparison results of the source and destination databases are consistent. Locate a comparison item you want to view and click **View Details** in the **Operation** column.

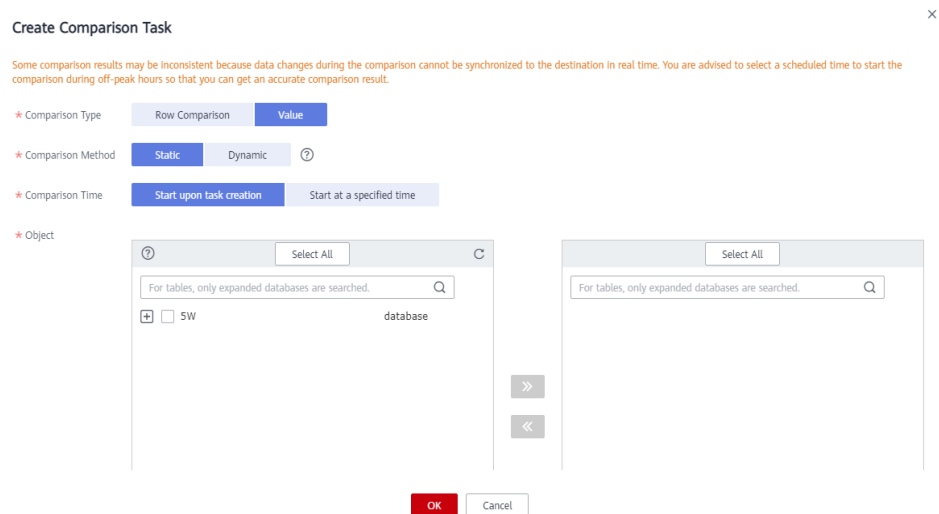
Figure 6-18 Object-level comparison



Item	Source Database	Destination Database	Result	Operation
Database	1	1	Consistent	View Details
Table	2	2	Consistent	View Details
Constraint	1	1	Consistent	View Details
Index	2	2	Consistent	View Details
Permission	4	4	Consistent	View Details

- On the **Data-Level Comparison** tab, click **Create Comparison Task**. In the displayed dialog box, specify **Comparison Type**, **Comparison Method**, **Comparison Time**, and **Object**. Then, click **OK**.

Figure 6-19 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 Comparison Value

* Comparison Method: Static Dynamic

* Comparison Time: Start upon task creation Start at a specified time

* Object

For tables, only expanded databases are searched.

SW database

OK Cancel

- **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.

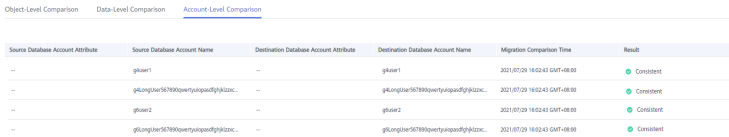
Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.
- **Comparison Policy:** DRS supports one-to-one and many-to-one comparison 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.
- **Object:** You can select objects to be compared based on the scenarios.
- **Account comparison:** Click the **Account-Level Comparison** tab to view the comparison results of database accounts and permissions.

Figure 6-20 Account-level comparison



Source Database Account Attribute	Source Database Account Name	Destination Database Account Attribute	Destination Database Account Name	Migration Comparison Time	Result
--	pkuser1	--	pkuser1	2021-07-29 16:52:43 GMT+08:00	Consistent
--	g@longuser567890qwertyuiopasdfghjklzxc...	--	g@longuser567890qwertyuiopasdfghjklzxc...	2021-07-29 16:52:43 GMT+08:00	Consistent
--	pkuser2	--	pkuser2	2021-07-29 16:52:43 GMT+08:00	Consistent
--	g@longuser567890qwertyuiopasdfghjklzxc...	--	g@longuser567890qwertyuiopasdfghjklzxc...	2021-07-29 16:52:43 GMT+08:00	Consistent

 NOTE

- Full synchronization tasks do not support account comparisons.
- Only PostgreSQL to PostgreSQL synchronization supports account comparison.


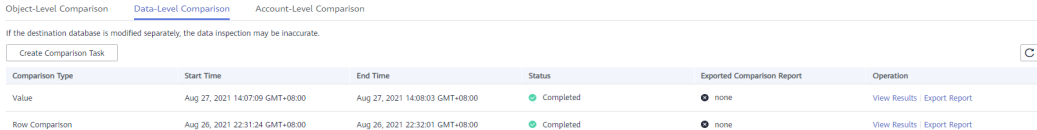
Step 4 After the comparison creation task is submitted, the **Data-Level Comparison** tab is displayed. Click  to refresh the list and view the comparison result of the specified comparison type.

Figure 6-21 Data-level comparison



Comparison Type	Start Time	End Time	Status	Exported Comparison Report	Operation
Value	Aug 27, 2021 14:07:59 GMT+08:00	Aug 27, 2021 14:08:03 GMT+08:00	Completed	none	View Results Export Report
Row Comparison	Aug 26, 2021 22:31:24 GMT+08:00	Aug 26, 2021 22:32:01 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 6-22 Viewing data-level comparison details

Comparison type: Row Comparison Comparison start time: Aug 26, 2021 22:31:34 GMT+08:00 Comparison end time: Aug 26, 2021 22:32:01 GMT+08:00

Source Database	Destination Database	Result	Operation
Testdb5	Testdb5	Consistent	View Details
factant	factant	Consistent	View Details
mgo	mgo	Consistent	View Details
testSpecMysql@192.168.1.10_+=[0,1]	testSpecMysql@192.168.1.10_+=[0,1]	Consistent	View Details
testdb1	testdb1	Consistent	View Details
testdb3	testdb3	Consistent	View Details
yjob	yjob	Consistent	View Details

Details Testdb5 - Testdb5

Source Database Table Name	Destination Database Table Name	Source Database Table Rows	Destination Database Table Rows	Row	Difference
col1	col1	47	47	Consistent	0
col1	col1	47	47	Consistent	0
colx	colx	28	28	Consistent	0

NOTE

You can also view comparison details of canceled comparison tasks.

----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

- Currently, only MySQL to GaussDB(DWS) synchronization supports periodic comparison.
- Many-to-one tasks do not 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 6-23 Modify Comparison Policy

×

Modify Comparison Policy

Status View comparison results in Data-Level Comparison.

Comparison Frequency Weekly ▾
A high comparison frequency may affect your service performance. Set a proper frequency based on service requirements.

Comparison Time Monday
 Tuesday
 Wednesday
 Thursday
 Friday
 Saturday
 Sunday

Time Zone GMT+08:00

Effective Time : 00–
 : 00
Periodic comparisons performed during off-peak hours have minor impacts on service performance and provides accurate comparison results. Comparisons that are not completed within the effective time will be automatically

----End

6.5 Managing Objects

6.5.1 Editing Synchronization Objects

This section describes how to change synchronization objects in a synchronization task. After a data synchronization task is created, you can change synchronization objects by adding or deleting databases and tables to be synchronized during the incremental synchronization.

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 6-24 Editing synchronization objects

<input type="checkbox"/>	Task Name/ID	Status	Data Fl...	DB E...	Synchroni...	Created	Netw...	Enter...	Operation
<input type="checkbox"/>		Incremental	To the c...	MySQL	Full+Incre...	Jan 06, 2022 18:...	VPN ...	default	Edit Stop More

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 the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.
- The name of the selected synchronization object cannot contain spaces.

Step 3 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 [Checking Whether the Source Database Is Connected](#) in *Data Replication Service User Guide*.

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

Figure 6-25 Pre-check

Check Again

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 4 In the synchronization task list on the **Data Synchronization Management** page, the current task status is **Incremental synchronization**, and a subtask in the

Modifying task status is generated. 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 **Step 4** from method 1.

----End

6.5.2 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 format of the object information in the template must meet the requirements. The value is case-sensitive and cannot include angle brackets (<>), periods (.), and double quotation marks ("). Objects that start or end with a space 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 6-26 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

6.5.3 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. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. The object types that can be mapped include database, schema, and table.

Object name mapping can be used only in the following scenarios:

- For the first time you select synchronization objects for a data synchronization task.
- For the first time you add or delete the synchronization object which is not in a mapping relationship.

NOTE

- Objects whose database names or table names contain newline characters cannot be mapped.
- If the destination DB is a type of PostgreSQL (for example, GaussDB, or PostgreSQL), the destination schema name cannot start with **pg_**. Otherwise, the migration fails.
- 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

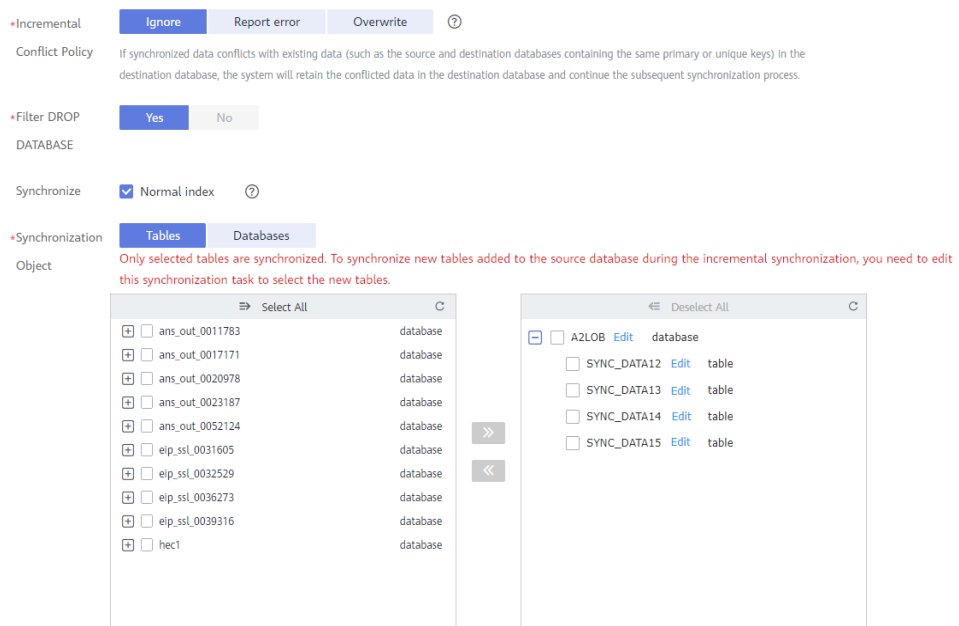
This section describes how to map objects when configuring a data synchronization task. For details about the mapping relationship, see [Viewing Synchronization Mapping Information](#).

Mapping Databases

During real-time synchronization, if the names of source databases to be synchronized are different from those in the destination, you can map the source database names to the destination ones.

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**.

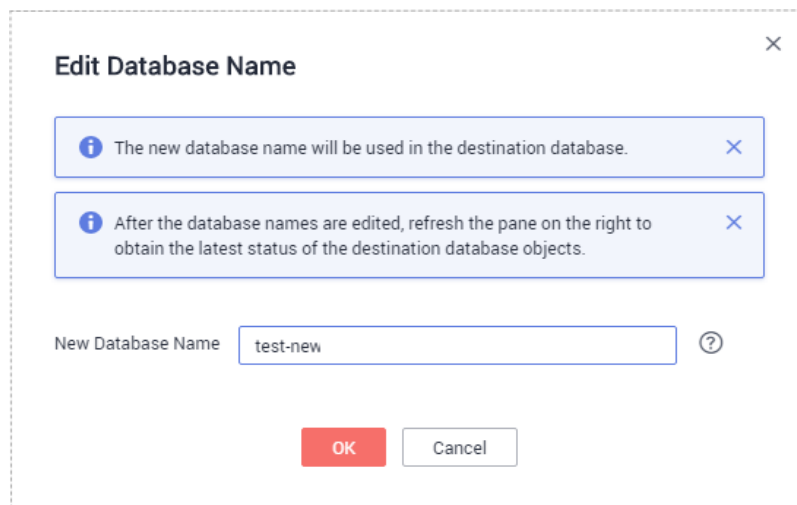
Figure 6-27 Mapping databases



Step 2 Changing a database name

In the displayed dialog box, enter a new database name. The new database name is the name of the database saved in the destination DB instance.

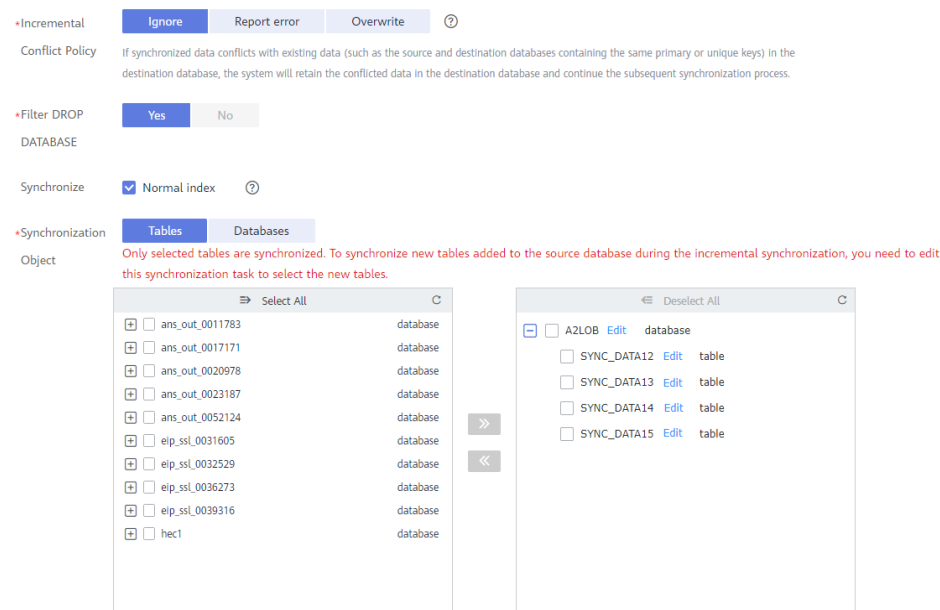
Figure 6-28 Changing a database name



Step 3 Check the result.

After the database name is changed, the database name before modification and the new database name are displayed. The database mapping is complete.

Figure 6-29 Checking the result



----End

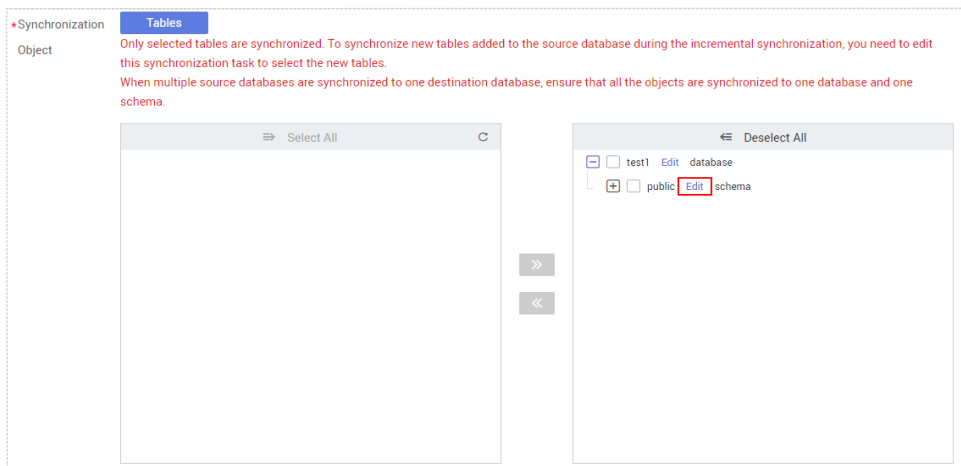
Mapping Schemas

A schema is a collection of database objects, including: tables, views, stored procedures, and indexes.

During real-time synchronization, if the names of source schemas to be synchronized are different from those in the destination, you can map the source schema names to the destination ones.

Step 1 On the **Set Synchronization Task** page, select the schema that needs to be mapped from the synchronization objects on the right area and click **Edit**.

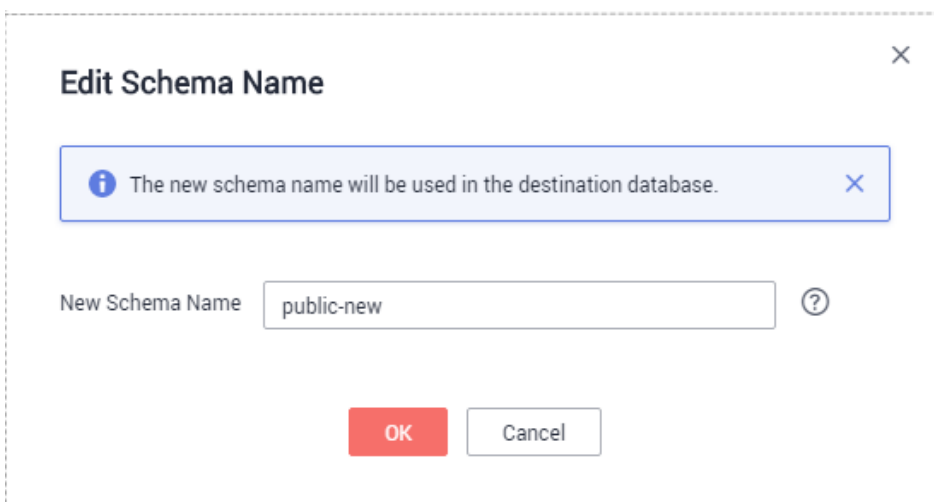
Figure 6-30 Mapping schemas



Step 2 Edit the schema name.

In the displayed dialog box, enter a new schema name which is the name to be saved in the destination database.

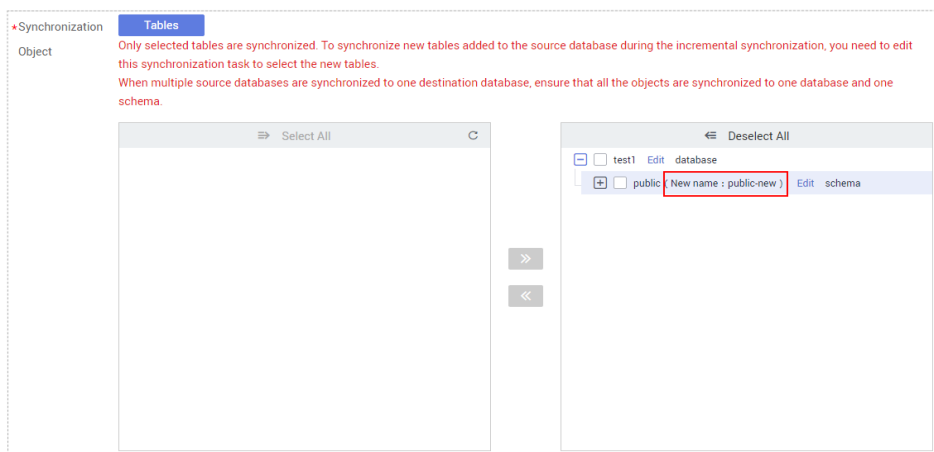
Figure 6-31 Edit Schema Name



Step 3 Check the result.

After the schema name is changed, the schema name before modification and the new schema name are displayed. The schema mapping is complete.

Figure 6-32 Checking the result



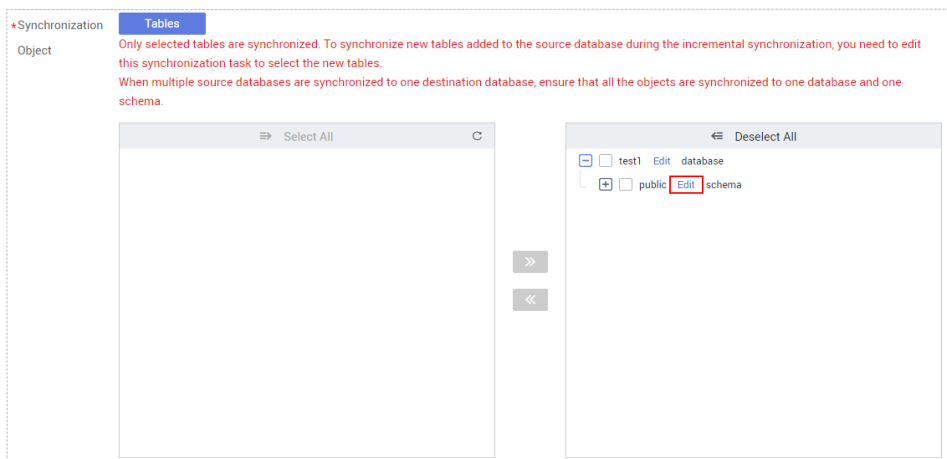
----End

Mapping Tables

During real-time synchronization, if the names of source tables to be synchronized are different from those in the destination, you can map the source table names to the destination ones.

- Step 1** On the **Set Synchronization Task** page, select the table that needs to be mapped from the synchronization objects on the right area and click **Edit**.

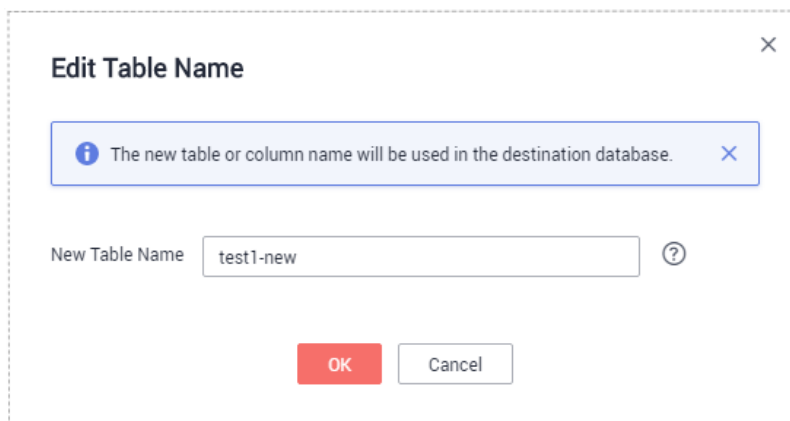
Figure 6-33 Table name mapping



- Step 2** Change a table name.

In the displayed dialog box, enter a new table name. The new table name is the name of the table saved in the destination database.

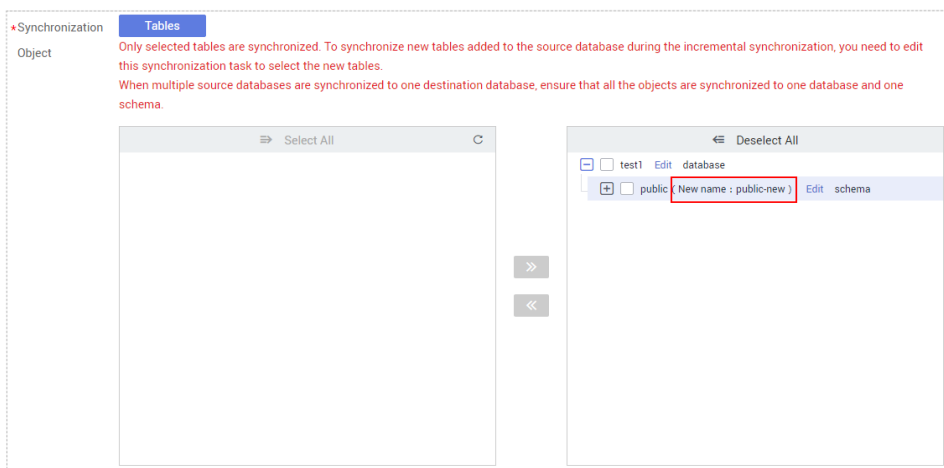
Figure 6-34 Changing a table name



Step 3 Check the result.

After the table name is changed, the table name before modification and the new table name are displayed. The table mapping is complete.

Figure 6-35 Checking the result



----End

6.5.4 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 6-36 Viewing mapping information

Source Database			Destination Database			
Database Name	Schema Name	Table Name	Database Name	Schema Name	Table Name	Operation
testdb330_1	sch1	test_char	testdb330_1	sch1	test_char	Edit
testdb330_1	sch1	test_int	testdb330_1	sch1	test_int	Edit View Column Processing

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

6.5.5 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 6-10](#).

Table 6-10 Data flow types that support data processing

Sync roni zatio n Direc tion	Data Flow	Data Filtering	Addition al Column	Column Processin g
To the cloud	MySQL->MySQL	Supporte d	Supporte d	Supporte d
To the cloud	MySQL->GaussDB(DWS)	Supporte d	Supporte d	Not supported
To the cloud	MySQL->	Supporte d	Supporte d	Not supported
To the cloud	DDM->GaussDB(DWS)	Not supported	Supporte d	Not supported

Sync hroni zatio n Direc tion	Data Flow	Data Filtering	Additional Column	Column Processin g
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
To the cloud	Oracle -> GaussDB distributed	Supporte d	Not supported	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	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 -> GaussDB(DWS)	Supporte d	Not supported	Not supported
From the cloud	GaussDB primary/standby -> GaussDB distributed	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 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 -> GaussDB distributed	Supporte d	Not supported	Supporte d
From the cloud	GaussDB distributed->GaussDB primary/standby	Supporte d	Not supported	Supporte d
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
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

Adding Synchronization Timestamp

Step 1 On the **Process Data** page of the real-time synchronization task, select the table objects to be processed in the **Select Source Database Object**.

Step 2 Click the **Process Data** area, click **Add Rule**, and enter the column name, type, and operation to be added.

NOTE

- The column to which the rule is to be added already exists in the table and cannot be the primary key.
- You are advised to use columns whose data type is timestamp (TIMESTAMP) as rule columns.

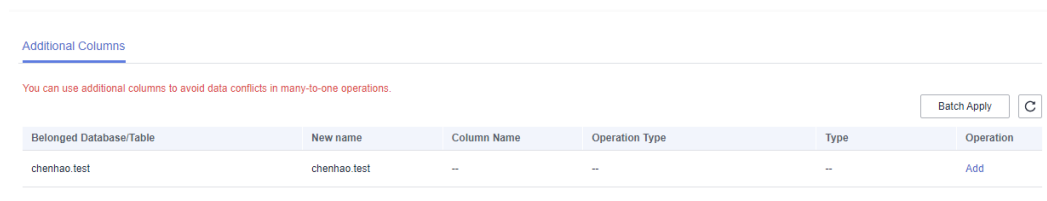
Step 3 Click **Next**.

----End

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.

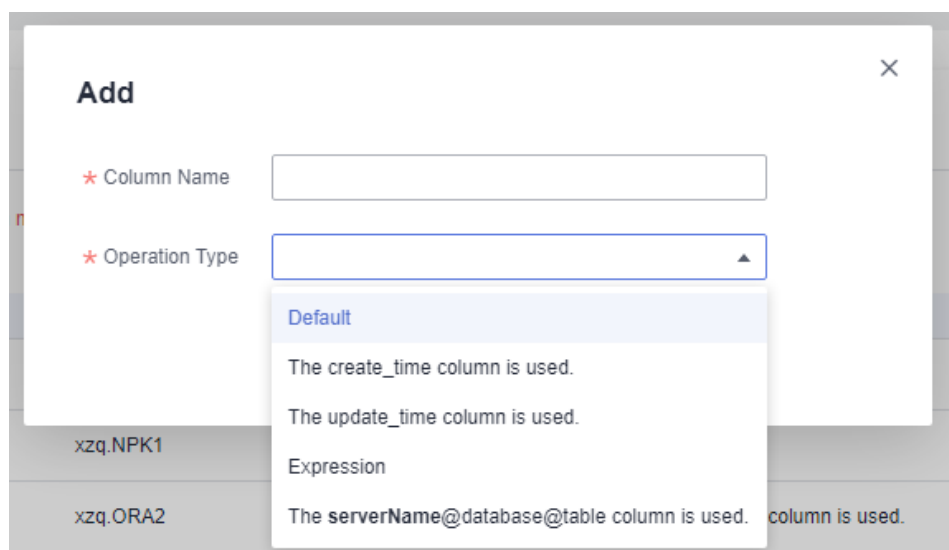
Figure 6-37 Additional columns



Belonged Database/Table	New name	Column Name	Operation Type	Type	Operation
chenhao.test	chenhao.test	--	--	--	Add

Step 2 In the displayed **Add** dialog box, specify the column name, operation type, and field type. Click **OK**.

Figure 6-38 Operation types



Add ✕

* Column Name

* Operation Type

Default

The create_time column is used.

The update_time column is used.

Expression

The serverName@database@table column is used. column is used.

xzq.NPK1

xzq.ORA2

 **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:** indicates that you need to fill the new column with 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.

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 6-39 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 500 tables can be filtered at a time.
- 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) must be enclosed in single quotation marks, for example, `update` > '2022-07-13 00:00:00' and age >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.
- 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.

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**.

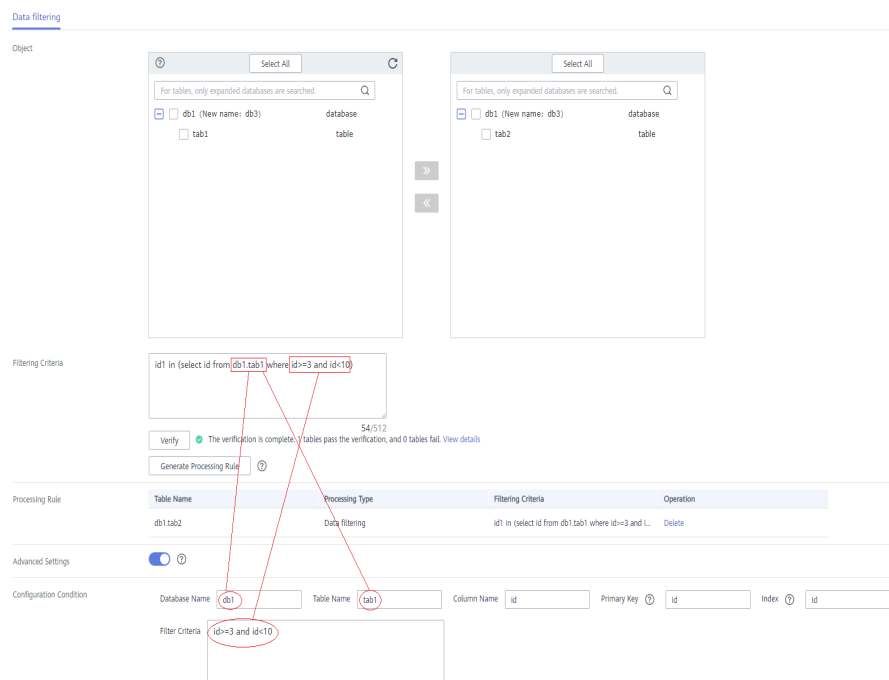
NOTE

- Each table has only one verification rule.
- Up to 500 tables can be filtered at a time.
- 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.
- 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.
- 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.

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 6-40 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**.

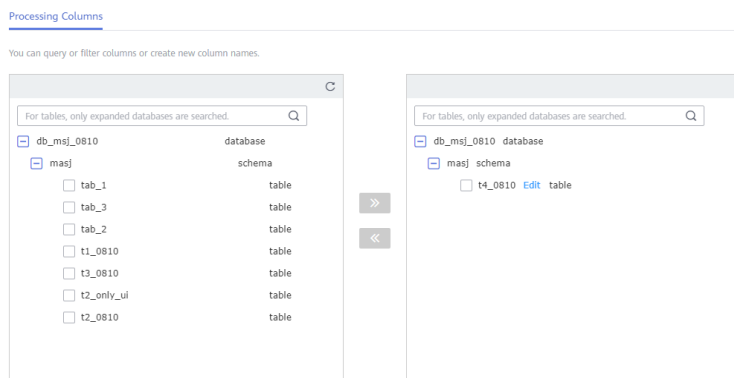
----End

Processing Columns

Step 1 On the **Process Data** page of the real-time synchronization task, select **Processing Columns**.

Step 2 In the **Object** area, select the objects to be processed.

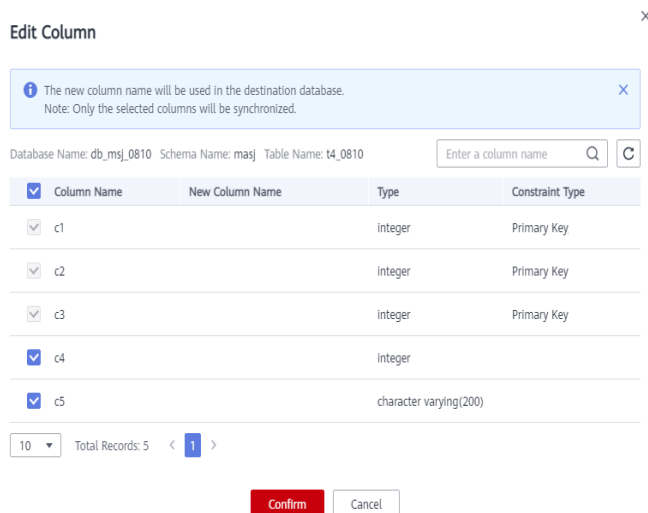
Figure 6-41 Processing Columns



Step 3 Click **Edit** to the right of the selected object.

Step 4 In the **Edit Column** dialog box, select the columns to be mapped and enter new column names.

Figure 6-42 Edit 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 the selected columns can be synchronized.
- MySQL to MySQL and MySQL to GaussDB(for MySQL) synchronizations do not support column mapping based on the partitioning column of a partitioned table.
- GaussDB partition keys cannot be filtered.


Step 5 Click **Confirm**.

Step 6 Click **Next**.

----End

Viewing Data Filtering Results

Step 1 In the task list, 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 task management page, click the target task name in the **Task Name/ID** column.


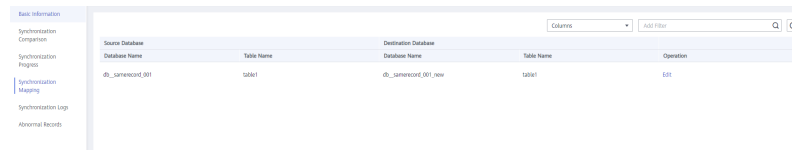
Step 2 In the navigation pane on the left, choose **Synchronization Mapping**. In the upper right corner, and select **Columns** to view column mapping records. Click  in the upper right corner to refresh the record list.

Figure 6-43 Viewing column mappings



The screenshot shows a web interface for 'Task Information' with a navigation menu on the left. The 'Synchronization Mapping' option is selected. The main area displays a table with columns for 'Source Database', 'Table Name', 'Destination Database', 'Table Name', and 'Operation'. A search bar at the top right is set to 'Columns'.

Source Database	Table Name	Destination Database	Table Name	Operation
db_sample001	tbl01	db_sample001_new	tbl01	EST

----End

6.6 Task Life Cycle

6.6.1 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

6.6.2 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 6-11 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.

- 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 6-44 Batch Operations



Step 3 In the displayed dialog box, enter the configuration information and click **Yes** to submit the configuration task.

----End

6.6.3 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:

- Source database password
- Destination database password

NOTE

After the preceding information is changed, the change takes effect immediately, and the data in the destination database is not cleared.

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, click **Modify Connection Details** in the **Connection Information** area.

Step 3 In the displayed dialog box, change the passwords of the source and destination databases and click **OK**.

Step 4 View the change result on the **Basic Information** tab.

----End

6.6.4 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.

NOTE

- 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** next to the **Flow Control** field.

Step 2 In the displayed dialog box, modify the settings.

Figure 6-45 Modifying the Flow Control Mode

Modify Flow Control

Flow Control Yes No ?

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

Method 2

Step 1 In the task list on the **Data Migration Management** page, locate the target task and choose **More > Speed** or **Speed** in the **Operation** column.

Figure 6-46 Task List

<input type="checkbox"/>	Task Name/ID <small>⌵</small>	Status	Charging	Data Flow	DB Engine <small>⌵</small>	Migration Type	Created <small>⌵</small>	Netwo...	Operation
<input type="checkbox"/>	DRS-4866	● Confi...	⊖ No	Out of th...	MySQL	Full+Incre...	Jan 07, 2022 10:14:27 ...	Public ...	Edit Stop Speed

Step 2 In the displayed dialog box, modify the settings.

Figure 6-47 Modifying the Flow Control Mode

Modify Flow Control

Flow Control Yes No ?

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.

Modify Flow Control

Flow Control Yes No ?

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->PostgreSQL
- MySQL->GaussDB(DWS)
- Oracle->MySQL
- Oracle->GaussDB(for MySQL)
- Oracle -> GaussDB(for MySQL) distributed
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- Oracle->GaussDB(DWS)
- Oracle->DDM
- Oracle->PostgreSQL
- PostgreSQL->PostgreSQL
- 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
- MongoDB->DDS

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- 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 -> Oracle
- GaussDB distributed -> GaussDB(DWS)
- GaussDB distributed -> GaussDB distributed
- GaussDB distributed->GaussDB primary/standby

Self-built -> Self-built

- MySQL->CSS/ES
- MySQL->Kafka
- Oracle->Kafka

- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed

6.6.5 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

6.6.6 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 left corner.


----End

Resume 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 6-48 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.6.7 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.

6.6.8 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

- After the task is paused, the status of the task becomes **Paused**.
- 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. It is recommended that the pause duration be less than or equal to 24 hours.


----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 6-49 Batch Operations



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(DWS)
 - PostgreSQL->PostgreSQL
 - PostgreSQL->GaussDB(DWS)
 - DDM->GaussDB(DWS)
 - Oracle->GaussDB(DWS)
 - Oracle->PostgreSQL
 - Oracle->GaussDB(for MySQL)
 - Oracle -> GaussDB primary/standby
 - Oracle -> GaussDB distributed
 - Oracle->RDS for MySQL
 - 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
 - MongoDB->DDS
- From the cloud
 - MySQL->MySQL
 - MySQL->CSS/ES
 - MySQL->Oracle
 - MySQL->Kafka
 - DDM->Oracle
 - DDM->Kafka
 - DDS->MongoDB
 - GaussDB(for MySQL)->MySQL
 - GaussDB(for MySQL)->Kafka
 - GaussDB(for MySQL)->CSS/ES

- GaussDB(for MySQL)->Oracle
- GaussDB primary/standby -> MySQL
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> Kafka
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - MySQL->CSS/ES
 - 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

In addition, the following tasks can be paused during full synchronization:

- MySQL->MySQL
- MySQL->GaussDB(for MySQL)
- GaussDB(for MySQL)->MySQL

6.6.9 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

- For a many-to-one synchronization task, only the parent task can be reset.
- For a MySQL many-to-one synchronization task, only the subtask can be reset.
- 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.

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 left 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(DWS)
- PostgreSQL->PostgreSQL
- DDM->DDM
- Oracle->MySQL
- Oracle->GaussDB(for MySQL)
- Oracle->GaussDB(DWS)
- Oracle->PostgreSQL
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- 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
- MongoDB->DDS

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- DDM->DDM
- DDS->MongoDB
- GaussDB(for MySQL)->Kafka
- GaussDB(for MySQL)->CSS/ES
- GaussDB(for MySQL)->Oracle
- GaussDB(for MySQL)->MySQL
- 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 -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->Kafka

Self-built -> Self-built

- MySQL->Kafka
- MySQL->CSS/ES
- 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

6.6.10 Skipping a Synchronization Task


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 6-50 Skip DDL



- Step 3** After confirming that statement was executed on the destination database, click **Yes** to skip the error and continue the synchronization task.

Figure 6-51 Skip DDL



----End

6.6.11 Performing a Primary/Standby Switchover

When you create a synchronization task, you can select the DRS instance type. The DRS instance can be a single node instance or a primary/standby instance. The primary/standby deployment can help you improve task reliability. If the task on the primary node fails, the standby node will take over the task and continue the synchronization.

You can select the DRS instance type in the following scenarios:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->GaussDB(DWS)
- From the cloud
 - MySQL->Kafka
 - GaussDB(for MySQL)->Kafka
 - GaussDB primary/standby -> Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - Oracle->Kafka
 - GaussDB primary/standby -> 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, and click **Next**.

Figure 6-52 Synchronization instance information




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 6-53 Primary/Standby tasks




Step 4 After a synchronization task is configured and started, DRS will start the task on the primary node, and the task on the standby node is suspended.

Figure 6-54 Before the primary/standby switchover



Step 5 If the task on the primary node is abnormal, DRS automatically starts the task on the standby node to continue the synchronization.

Figure 6-55 After the primary/standby switchover



----End

6.6.12 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
 - PostgreSQL->PostgreSQL
- From the cloud
 - MySQL->MySQL
 - MySQL->CSS/ES
 - 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


6.6.13 Changing Specifications

You can change the DRS instance 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. Only whitelisted users can use this function. You need to submit a service ticket to apply for this function.

Constraints

- You can change the task specifications only when your account balance is more than \$0 USD.
- Currently, DRS supports specification upgrade only in MySQL-to-MYSQL synchronization tasks with single-node DRS instances configured. Task specifications cannot be downgraded.
- 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 instance specifications 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 6-56 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 [jacecaa5-782c-4a9b-9344-d5cb221a201] completed, total item 4, success item 4, not pass item 0
2022/12/02 19:14:32 GMT+08:00	Info	precheck [jacecaa5-782c-4a9b-9344-d5cb221a201] start
2022/12/02 19:14:18 GMT+08:00	Info	precheck [jacecaa5-782c-4a9b-9344-d5cb221a201] completed, total item 4, success item 4, not pass item 0
2022/12/02 19:13:57 GMT+08:00	Info	precheck [jacecaa5-782c-4a9b-9344-d5cb221a201] 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, caused 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

6.6.14 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.


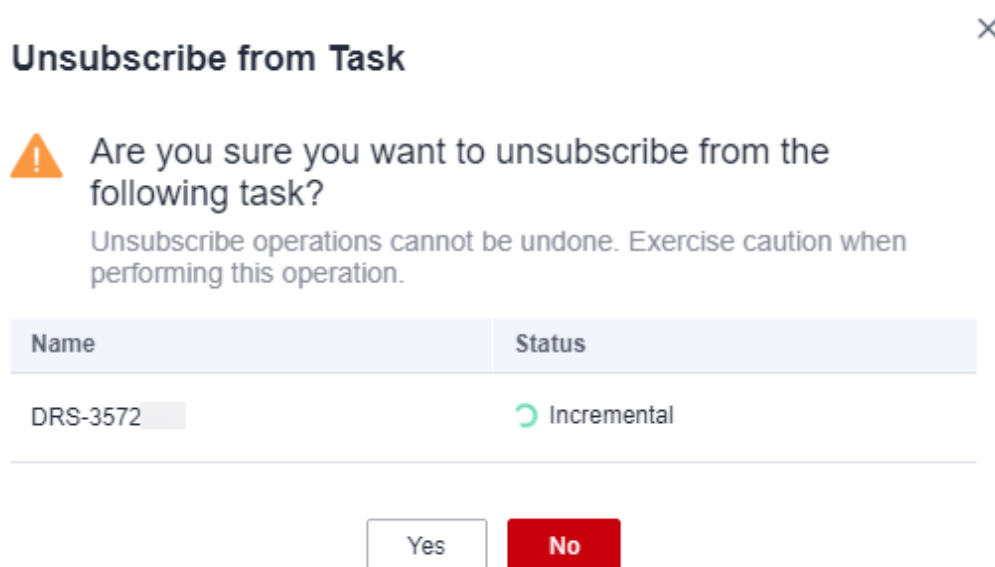
- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** On the **Data Synchronization Management** page, select the 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.

Figure 6-57 Unsubscribing from a task



- 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** 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.
 - Alternatively, search for target orders by name, order No., or ID in the search box.
- Step 6** On the displayed page, confirm the order to be unsubscribed from and select a reason. Then, click **Confirm**.
- Step 7** 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

6.6.15 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 **OK**.

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**.


----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 6-58 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.6.16 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 6-59 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.6.17 Task Statuses

Synchronization statuses indicate different synchronization phases.

Table 6-12 lists synchronization task statuses and descriptions.

Table 6-12 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.

Status	Description
Starting	The task is being started.
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.
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 ends.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.
- Deleted synchronization tasks are not displayed in the status list.

7 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 10 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 On the **Basic Information** tab, click the **Tags** tab.

Step 3 On the **Tags** tab, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.

Add Tag ✕

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#)

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

- 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 After a tag has been added, you can view and manage it 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 On the **Basic Information** tab, click the **Tags** tab.

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 On the **Basic Information** tab, click the **Tags** tab.

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**

8 Diagnosis of Abnormal Records

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**

Currently, only the following data flows support this function: DDM->GaussDB(DWS), DDM->Oracle, Oracle->GaussDB(DWS), Oracle->RDS for MySQL, Oracle->GaussDB(for MySQL), Oracle->PostgreSQL, MySQL->GaussDB(DWS), MySQL->GaussDB, and MySQL->PostgreSQL.

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

9 Interconnecting with CTS

9.1 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 9-1 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

9.2 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

10 Interconnecting with Cloud Eye

10.1 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 10-1](#) lists the DRS performance metrics.

Table 10-1 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 bytes/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 bytes/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 bytes/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 bytes/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 bytes/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.	≥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 bytes/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	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
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

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
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
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

Dimensions

Key	Value
instance_id	DRS instance ID

10.2 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**.

Figure 10-1 Choosing a monitored object



- 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.

Figure 10-2 Configuring alarm information

* Name

Description

* Enterprise Project [Create Enterprise Project](#)

* Resource Type Data Replication Service

* Dimension DRS

* Monitoring Scope Specific resources

* Monitored Object DRS

* Method Use template Create manually

* Template [Create Custom Template](#)

Alarm Policy	Alarm Severity	Operation
Trigger an alarm if CPU Usage Raw data >= 23% for 3 consecutive periods. Trigger an alarm one day again if the alarm persists.	Major	Delete

Alarm Notification

* Notification Recipient Notification group Topic subscription

- 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

10.3 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 10-3 Viewing monitoring metrics



----End

11 Operation Reference in Synchronization Scenarios

11.1 Kafka Message Format

Data synchronized to the Kafka cluster is stored in Avro, JSON, and JSON-C formats.

Avro Format

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. For details about the data parsing, see [drs-avro-read.rar](#).

JSON

For details about the JSON format from MySQL and GaussDB(MySQL) to Kafka, see [Table 11-1](#). For details about the JSON format from GaussDB, PostgreSQL, and Oracle to Kafka, see [Table 11-2](#).

Table 11-1 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.

Parameter	Description
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

```
{
  "mysqlType":{
    "c1":"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": "asfiajhfi939-0239uoituqorjoqirfoidfqrniowejoiwqjroqwjrowqjoioiqgoiegnkjgoi23roiugou
      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": "asfiajhfi939-0239",
      "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 11-2 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.
dbType	Source database type The types of different engines are as follows: GaussDB Primary/Standby GaussDB Distributed PostgreSQL Oracle
schema	Name of a scheme
opType	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
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 previous transaction. The value is a 13-digit Unix timestamp, in milliseconds. GaussDB distributed: commit time of the previous 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.
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.

Parameter	Description
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 11-3](#).

Table 11-3 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 .

Parameter	Description
pkNames	Primary key name

Common Escape Characters in JSON

Table 11-4 Escape Character

Character	Escape character
<	\u003d
>	\u003e
&	\u0026amp;
=	\u003d
'	\u0027

11.2 Kafka Authentication

PLAINTEXT

No security authentication mode is available. You only need to enter the IP address and port for connection.

Figure 11-1 PLAINTEXT

Destination Database

IP Address or Domain Name



Ensure that the entered addresses belong to the same DB instance.

Security Protocol



Test Connection

SASL_PLAINTEXT

The SASL mechanism is used to connect to Kafka, and you need to configure SASL parameters.

Figure 11-2 SASL_PLAINTEXT

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

SASL Mechanisms ?

Username

Password

Table 11-5 Parameter settings

Parameter	Description
SASL Mechanisms	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 . <ul style="list-style-type: none"> ● GSSAPI ● PLAIN ● SCRAM-SHA-256 ● SCRAM-SHA-512
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 11-3 SSL

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

Truststore Certificate

Truststore Certificate Password ?

Endpoint Identification Algorithm ?

Mutual SSL Authentication

Keystore Certificate

Truststore Certificate Password ?

Keystore Private Key Password ?

Table 11-6 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
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 11-4 SASL_SSL

Destination Database

IP Address or Domain Name	<input type="text"/>	?	
Ensure that the entered addresses belong to the same DB instance.			
Security Protocol	SASL_SSL	?	
SASL Mechanisms	GSSAPI	?	
Username	<input type="text"/>		
Password	<input type="text"/>		
Truststore Certificate	<input type="text"/>	Select	
Truststore Certificate Password	<input type="text"/>	?	
Endpoint Identification Algorithm	<input type="text"/>	?	
Mutual SSL Authentication	<input checked="" type="checkbox"/>		
Keystore Certificate	<input type="text"/>	Select	
Truststore Certificate Password	<input type="text"/>	?	
Keystore Private Key Password	<input checked="" type="checkbox"/>	<input type="text"/>	?
<input type="button" value="Test Connection"/>			

11.3 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.

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

11.4 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.

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

11.5 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.

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.relnamespace)||','||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

11.6 Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL

This section describes how to perform real-time synchronization from PostgreSQL to RDS PostgreSQL. You can create 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, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME)
 - 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, and ALTER 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

 **NOTE**

If the source is an RDS for PostgreSQL instance on the current cloud, you can create related objects as user **root**. If the "Must be superuser to create an event trigger" error is reported, you can submit a service ticket. For details about permissions of user **root** of RDS for PostgreSQL on the current cloud, see Relational Database Service User Guide.

Step 1 Connect to the database to be synchronized as a user who has permission to create event triggers.

Step 2 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_address  varchar(64) default inet_client_addr(),
  client_port integer default inet_client_port(),
  event_time timestamp default current_timestamp
);
```

Step 3 Run the following statements to create a 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')) 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;
$BODY$;
```

Step 4 Run the following statements to grant necessary permissions to the objects created in [Step 2](#) and [Step 3](#):

```
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;
```

Step 5 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();
```

Step 6 Run the following statement to set the created event trigger to enable:

```
ALTER EVENT TRIGGER hwdrs_ddl_event ENABLE ALWAYS;
```

Step 7 Return to the DRS console and create a PostgreSQL to RDS PostgreSQL synchronization task.

Step 8 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;
```

----End

A Change History

Released On	Description
2022-10-20	This issue is the second official release, which incorporates the following change: <ul style="list-style-type: none">When the source database is an Oracle database, the WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13 and WE8ISO8859P15 character sets are supported.
2022-09-30	This issue is the first official release.